

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
6 September 2002 (06.09.2002)

PCT

(10) International Publication Number
WO 02/069608 A2

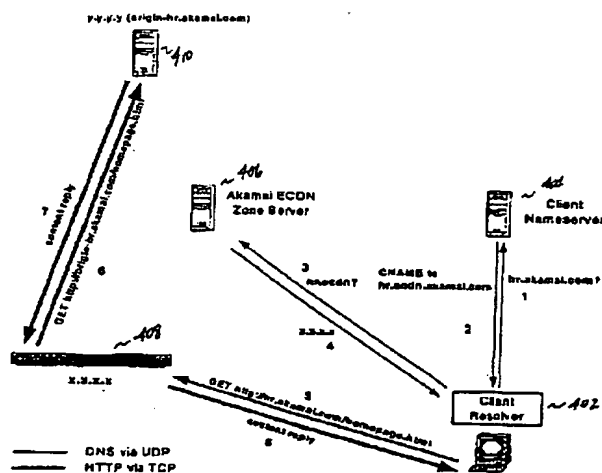
- (51) International Patent Classification⁷: **H04L 29/12, 29/06** (74) Agent: JUDSON, David, H.; Akamai Technologies, Inc., 500 Technology Square, Cambridge, MA 02139 (US).
- (21) International Application Number: **PCT/US02/01079** (81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZM, ZW.
- (22) International Filing Date: 16 January 2002 (16.01.2002)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
60/262,171 16 January 2001 (16.01.2001) US
- (71) Applicant (*for all designated States except US*): AKAMAI TECHNOLOGIES, INC. [US/US]; 500 Technology Square, Cambridge, MA 02139 (US).
- (72) Inventor; and
- (75) Inventor/Applicant (*for US only*): NEERDAELS, Charles, Joseph [US/US]; 622 Bayview Drive, Aptos, CA 95003 (US).
- (84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declaration under Rule 4.17:

— of inventorship (Rule 4.17(iv)) for US only

[Continued on next page]

(54) Title: USING VIRTUAL DOMAIN NAME SERVICE (DNS) ZONES FOR ENTERPRISE CONTENT DELIVERY



- 1) Client resolver makes DNS request to its name server (NS)
2) Local NS looks up hr.akamai.com, and so it returns a CNAME redirect to ecdn.akamai.com (the shorter name was simply a matter of convenience)
3) Client resolver makes DNS request to ecdn zone NS, for hostname hr.ecdn
4) ECDN zone server replies with an address x.x.x.x, or could optionally pass it to another ECDN zone server for hr, or could CNAME to another hostname
5) Assuming it resolves, the client makes an http request to x.x.x.x for with host hr.akamai.com
6,7,8,9) ECDN is configured to be a surrogate for origin-hr.akamai.com, and satisfies the HTTP request accordingly

(57) Abstract: According to the invention, a domain to be published to an enterprise ECDN is associated (either by static configuration or dynamically) with a set of one or more enterprise zones configurable in a hierarchy. When a DNS query arrives for a hostname known to be associated with given content within the control of the ECDN, a DNS server preferably responds in one of three (3) ways; handing back an IP address, e.g., for an ECDN intelligent node that knows how to obtain the requested content from a surrogate or origin server; executing a zone referral to a next (lower) level name server in a zone hierarchy, or CNAMing to another hostname, thereby essentially restarting the lookup procedure. In the latter case, this new CNAME causes the resolution process to start back at the root and resolve a new path, probably along a different path in the hierarchy. At any particular level in the zone hierarchy, preferably there is an associated zone server. That server preferably executes logic that applied the requested hostname against a map, which, using known techniques, may be generated from given (static, dynamic, internally-generated or third party-sources) performance metrics. Thus, a given name query to ECDN-managed content may be serviced in coordination with various sources of distributed network intelligence. As a result, the invention provides for a distributed, dynamic globally load balanced name service.

WO 02/069608 A2

Best Available Copy



Published:

— without international search report and to be republished
upon receipt of that report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

USING VIRTUAL DOMAIN NAME SERVICE (DNS) ZONES FOR ENTERPRISE CONTENT DELIVERY

BACKGROUND OF THE INVENTION

This application is based on and claims priority from Provisional

5 Application Serial No. 60/262,171 filed January 16, 2001.

Technical Field

The present invention relates generally to content delivery and management within a private enterprise network.

Description of the Related Art

10 It is well-known to deliver digital content (e.g., HTTP content, streaming media and applications) using an Internet content delivery network (ICDN). A content delivery network or "CDN" is a network of geographically distributed content delivery nodes that are arranged for efficient delivery of content on behalf of third party content providers. A request from a requesting end user for given
15 content is directed to a "best" replica, where "best" usually means that the item is served to the client quickly compared to the time it would take to fetch it from the content provider origin server.

Typically, a CDN is implemented as a combination of a content delivery infrastructure, a request-routing mechanism, and a distribution infrastructure. The
20 content delivery infrastructure usually comprises a set of "surrogate" origin servers that are located at strategic locations (e.g., Internet network access points, Internet Points of Presence, and the like) for delivering copies of content to requesting end users. The request-routing mechanism allocates servers in the content delivery infrastructure to requesting clients in a way that, for web content delivery,
25 minimizes a given client's response time and, for streaming media delivery, provides for the highest quality. The distribution infrastructure consists of on-demand or push-based mechanisms that move content from the origin server to the surrogates. An effective CDN serves frequently-accessed content from a surrogate that is optimal for a given requesting client. In a typical CDN, a single service

provider operates the request-routers, the surrogates, and the content distributors. In addition, that service provider establishes business relationships with content publishers and acts on behalf of their origin server sites to provide a distributed delivery system. A well-known commercial CDN service that provides web
5 content and media streaming is provided by Akamai Technologies, Inc. of Cambridge, Massachusetts.

Enterprises have begun to explore the desirability of implementing content delivery infrastructures to address several problems. Currently, enterprise users typically experience slow and expensive access to Internet content. Slow access to
10 business critical data available on the Internet hurts productivity, and the cost of providing good access, e.g., by building bigger networks and by deploying and managing caching infrastructure, is large. In addition, many IT organizations cannot deliver the required quality of service for Internet content delivery due to lack of talent and expertise. Yet another reason corporations are exploring CDNs
15 is because of the slow, expensive and often cumbersome access to and within the entity's intranet. As corporate intranets quickly become a critical component of business process in many large companies, fast and efficient access to the data and applications on the intranet is a high priority for many IT departments. Nevertheless, current intranet delivery solutions are inadequate, and solving the
20 problems, e.g., by building bigger internal networks, deploying and managing caches, and distributing application front ends, is extremely expensive. To address these deficiencies, several large software vendors are attempting to build ecosystems to provide web-based front ends to many enterprise applications, however, distributing these application from front ends efficiently, in of itself, will
25 be a critical IT problem that current technologies do not address. Finally, enterprises are considering CDN technology due to slow, expensive access to business partner applications and information provided by current techniques and solutions. Business-to-business applications (such as ordering, inventory, and pricing management) between business partners is done today by linking partners

with a physical network. These applications are moving to the Internet/intranet, and the need to link business partners together in an efficient way with web-based front ends is another critical IT problem that is not addressed by today's solutions.

BRIEF SUMMARY OF THE INVENTION

5 It is a general object of the invention to define and implement one or more virtual zones within an enterprise namespace to facilitate enterprise content delivery behind a corporate firewall.

 It is another general object of the present invention to provide an enterprise content delivery network that coexists with existing or "legacy" Domain Name
10 Service (DNS) infrastructure to facilitate mapping of requests for enterprise resources to surrogate servers, e.g., common caching appliances, application servers, distributed storage and database management systems, and streaming media servers.

 It is yet another general object of the invention to define and implement
15 one or more so-called "virtual" zones within an enterprise namespace to facilitate content delivery behind a corporate firewall over an enterprise content delivery network (ECDN).

 Another object of the present invention to enable an ECDN DNS to coexist with an existing enterprise DNS and to enable content delivery with only minimal
20 configuration changes to the existing infrastructure using virtual zones. A given host or sub-domain to be published to the ECDN is aliased into a virtual zone namespace preferably using a DNS Canonical Name (CNAME).

 A still further object of the invention is to implement an enterprise content delivery network wherein any arbitrary hierarchical namespace can be
25 implemented and wherein each layer of the namespace inherits the namespace above.

 According to the invention, a domain to be published to an enterprise ECDN is associated (either by static configuration or dynamically) with a set of one or more enterprise zones configurable in a hierarchy. When a DNS query

arrives for a hostname known to be associated with given content within the control of the ECDN, a DNS server preferably responds in one of three (3) ways: (a) handing back an IP address, e.g., for an ECDN intelligent node that knows how to obtain the requested content from a surrogate or origin server; (b) executing a zone referral to a next (lower) level name server in a zone hierarchy, or (c) CNAMing to another hostname, thereby essentially restarting the lookup procedure. In the latter case, this new CNAME causes the resolution process to start back at the root and resolve a new path, probably along a different path in the hierarchy. At any particular level in the zone hierarchy, preferably there is an associated zone server. That server preferably executes logic that applies the requested hostname against a map, which, using known techniques, may be generated from given (static, dynamic, internally-generated or third party-sourced) performance metrics. Thus, a given name query to ECDN-managed content may be serviced in coordination with various sources of distributed network intelligence. As a result, the invention provides for a distributed, dynamic globally load balanced name service.

According to a more specific aspect of the invention, a domain to be published to the ECDN is associated (either by static configuration or dynamically) with a set of one or more enterprise zones configurable in a hierarchy. When a request for content is received from a given client in the enterprise (or otherwise accessible through the firewall), a name query to the enterprise's legacy DNS is directed to the configured ECDN zone primary server. This primary ECDN zone server associates, and relays via CNAME, a name query with an ECDN intelligent node (or resolves directly to a given content origin server). This association may be based on given conditions within the network, server load conditions, or some combination of various known metrics. With ECDN virtual zones, maps that associate a given request to a server may be local or global, static or dynamic. While the maps typically are generated within the ECDN infrastructure, they could also be affected by external agents providing

hints or modifications. If multiple zones are used in a hierarchy, a first level ECDN zone server directs the name query to a second level ECDN zone server, and so on, until the appropriate response is generated. At any one level the particular zone server need not be able to resolve the entire name; rather, the particular zone server need only be able to resolve the portion required for directing the next level of resolution. For example: the zone server for the namespace containing b in the hostname a.b.c.d.e may have no idea where the final resolution will end up, but it should know which set of servers may be authoritative for the namespace containing c. Because DNS readily allows one to "carve out" zones in the namespace of a particular domain (e.g., ecdn.company.com can be managed by a different authoritative DNS server than the rest of company.com), the creation of the first level ECDN zone on the existing DNS infrastructure allows the two systems to coexist and undergo modification in parallel without interference with one another. This enables the present invention to be readily integrated into an existing or legacy DNS solution with minimal reconfiguration.

Thus, according to the invention, an existing DNS infrastructure is augmented with one or more ECDN zone servers, and requests for content to be delivered over the ECDN are resolved through the one or more zone servers. In one embodiment, this function is achieved by identifying a given enterprise domain to be published to the ECDN, and having a network administrator (or a person acting on the administrator's behalf) modify an existing DNS name record for that enterprise domain. In this embodiment, the administrator CNAMEs the enterprise domain to a domain identified by a so-called "virtual zone" that is managed on an ECDN zone server (and, thus, forms part of the ECDN). In particular, this process enables the enterprise domain to be managed (by the enterprise or a third party acting on the enterprise's behalf) as part of the ECDN, typically by the ECDN zone server on one of the intelligent nodes (or some other machine, such as a client name server). A similar process is repeated for every

new domain or subdomain being distributed by the ECDN. Any client downloading a URL referring to the enterprise domain then gets directed to the ECDN zone server, which directs the client to an ECDN intelligent node. The client passes a host header to the ECDN node to identify the actual host.

- 5 The foregoing has outlined some of the more pertinent features of the present invention. These features should be construed to be merely illustrative. Many other beneficial results can be attained by applying the disclosed invention in a different manner or by modifying the invention as will be described.

BRIEF DESCRIPTION OF THE DRAWINGS

- 10 Figure 1 is a block diagram of an enterprise in which the inventive enterprise content delivery network (ECDN) may be implemented;

Figure 2 is another enterprise environment in which the present invention may be implemented;

- 15 Figure 3 illustrates a known technique for handling a request for content in an enterprise using a legacy DNS infrastructure;

Figure 4 illustrates a technique for enterprise content delivery using virtual DNS zones according to the present invention;

Figure 5 illustrates a representative DNS namespace in which an enterprise virtual zone is created according to the present invention;

- 20 Figure 6 illustrates a zone hierarchy in which first and second level ECDN zone servers are used to resolve a name query to an enterprise domain that has published to the ECDN; and

Figure 7 illustrates a technique for dynamic zone creation through an agent-hinted CNAME according to the present invention; and

- 25 Figure 8 illustrates a technique for using virtual zones to enable an enterprise partner to access an enterprise ECDN to retrieve intranet content.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

The present invention is implemented within an "enterprise" or "enterprise environment." Referring to Figure 1, a representative enterprise environment in

which the present invention is implemented comprises at least one primary central office 100 (e.g., with central IT support), which is connected to one or more regional datacenters 102, with each datacenter connected to one or more remote branch offices 104 via private line, or via the public Internet (most likely with a virtual private network or "VPN"). The branch offices 104 may be fully (or partially) meshed amongst themselves. Figure 2 illustrates another enterprise computing environment that includes a central office 200 and a pair of remote offices 202. As shown in Figure 2, remote office 202a is connected to the central office 200 over a private line, which refers to a line not generally routable over the public Internet (e.g., frame relay, satellite link, microwave link, or the like), and remote office 202b is connected to the central office 200 over a virtual private network (VPN) typically over the public Internet, or in other known ways. A given remote office may be an office of an enterprise business partner if an appropriate business relationship between the enterprise and the partner exists.

Enterprise networks are generally based on one of two network types – point-to-point private/leased line, or VPN over public lines (often fully meshed between offices). From a topological view, actual enterprise networks may not operate in a hierarchical sense for connectivity. From a logical management view, however, the servers remain in a relatively weak hierarchy (as illustrated in either Figure 1 or 2) with more valuable and management intensive systems usually existing at more centralized data centers, and very few of these types of systems being deployed in remote branched offices. This architecture is usually a function of where information and management reside, which drives the cost of management for the overall IT network and IT services. Thus, for example, in the more complex enterprise such as shown in Figure 1, regional datacenters tend to have more critical information than branch offices, and the central office tends to have the most critical information. It is assumed for purposes of illustration that the enterprise hosts given internal content that it desires to have stored, cached and delivered to end users in the enterprise. Such internal enterprise content, or

ECDN content, generally is non publicly-available content (in whatever format) that an enterprise desires to make available to permitted users within the enterprise or within a third party partner of the enterprise. An end user typically operates a computer having an Internet browser. The authoritative content and application servers 106 (in Figure 1) and 206 (in Figure 2) being published into the CDN (each identified here as an "ECDN server") can exist in any location, although from a practical perspective they may often be somewhat centrally-located. Such servers are sometimes referred to as origin servers. The enterprise is also assumed to be IP based and to have a Domain Name Service (DNS) 108 (in Figure 1) and 208 (in Figure 2). From a security standpoint, the enterprise network manager roughly divides the world of the network into trusted and un-trusted, which usually corresponds to internal and external entities. Firewalls 110 (in Figure 1) and 210 (in Figure 2) are typically configured to allow most outbound traffic, with severe restrictions on inbound traffic. This allows for access to requested services without providing third parties a foothold for intrusion. More sophisticated systems usually create a security entity called a DMZ, which can be thought of as a set of two firewalls, with certain assets like email, DNS, web servers, etc. sitting between them. Each firewall has a different set of filtering rules, with the innermost generally allowing valid traffic from a host within the DMZ to enter the enterprise. Generally, no externally initiated traffic is allowed to pass both sets of filters.

For purposes of illustration, an enterprise content delivery network (ECDN) is provisioned into the enterprise network topology of Figure 2. The ECDN may be implemented in a standalone manner, or it may be part of a larger Internet CDN, in which case the Internet CDN service provider may manage the ECDN on behalf of the enterprise. Various techniques for extending an ICDN into an enterprise are described in copending application Serial No. xx/xxx,xxx, filed January 7, 2002, titled **EXTENDING AN INTERNET CONTENT DELIVERY NETWORK INTO AN ENTERPRISE**, and in copending

application Serial No. yy/yyy,yyy, filed January 7, 2002, titled **EXTENDING AN INTERNET CONTENT DELIVERY NETWORK INTO AN ENTERPRISE ENVIRONMENT BY LOCATING ICDN CONTENT SERVERS TOPOLOGICALLY NEAR AN ENTERPRISE FIREWALL**, which

5 applications are commonly-owned and which are incorporated herein by reference. Referring back to Figure 2, ECDN functionality is provided by locating one or more intelligent nodes (each referred to as an IN) that execute appropriate management, provisioning, monitoring and routing applications to facilitate the distribution, delivery and management of the private enterprise content delivery
10 network. Thus, for example, an IN 212a may be located in the central office 200 and each of the remote offices may include an IN box 212b as indicated. An IN box 212 is typically a computer, namely, a server, having at least one Pentium-class processor, an operating system (e.g., Linux, Windows NT, Windows 2000, or the like), and some amount of disk storage and system memory. Using the IN
15 boxes, the enterprise (or the CDNSP on its behalf) has the ability to map and load balance users inside the enterprise (perhaps as part of the global ICDN), to fetch content from inside the firewall, to collect log information, to deploy software, to distribute live streams, and to provide other ICDN functionality.

Thus, an illustrative enterprise CDN ("ECDN") is comprised of a variety
20 of origin and surrogate servers and a distributed set of intelligence nodes (IN). An origin server can be any server that is intended to respond to a user request for information and includes, for example, common caching appliances, application servers, distributed storage and database management systems, and streaming media servers. The intelligence nodes tie the entire set of distributed nodes and
25 services into a single, virtual service cloud. The ECDN intelligence nodes may have different configurations. Thus, for example, a first type (e.g., IN 212a in central office 200) may be sized (in terms of processing power and memory) for the largest datacenters, and such machines would be aggregation points for monitoring, logging, management and request routing. A second type of

intelligence node (e.g., IN 212b) may be sized for branch office deployment, and such servers would be sized to optimize local behavior. The nodes may be placed in a hierarchy with an IN of the first type acting as a parent. An intelligence node of the second type would typically include a caching engine.

5 Architecturally, an ECDN such as illustrated above is similar to an Internet content delivery network (ICDN) and comprises: HTTP, streaming and DNS server nodes deployed across a customer network. Generalizing, the ECDN provides the enterprise with the following: a namespace and an associated request routing mechanism (to control how to navigate the CDN), storage, caching and
10 streaming (to control how and where are objects stored and distributed), optionally an application environment (to facilitate application extensibility), management and provisioning (to control how the CDN is deployed and managed), and reporting/auditing (to provide visibility into the behavior). An ECDN such as described above provides an application, streaming and content delivery platform
15 so that enterprise content may be delivered and managed in a decentralized fashion while providing a tightly-integrated solution with a customer's storage, network, application and management solutions.

Enterprises today either outsource their Domain Name Service (DNS) or, as described above, run a DNS infrastructure (e.g., one or more DNS servers). A
20 DNS namespace has certain characteristics, as are well-known. Each unit of data in a DNS distributed database is indexed by a hostname. These names are simply leaves in a tree representation, which is often called a domain namespace. Each node in the tree is labeled with a simple name. The fully qualified hostname of any node in the tree is the sequence of labels on the path from that node to the
25 root. A zone is a subtree of the domain name space. Names at the leaves of the tree generally represent individual hosts, whose names resolve to one or more IP addresses. It has become common in instances of virtual hosting within ISPs to have multiple customers share a single machine, therefore having multiple hostnames resolve to a single machine or server farm. CDN technology to date has

additionally made it common for a single hostname to potentially resolve to any one of thousands of machines, though at any given time the set of possible hostnames is configured and fixed. The term zone usually relates to a set of hosts within a single authoritative server's namespace, whereas a single zone could include multiple sub-domains. Thus, e.g., if foo.com and eng.foo.com are both
5 handled by the same nameserver, they are both in the same zone; but, if foo.com's nameserver refers a lookup to a different authority for resolution, they would be considered separate zones.

An organization administrating a domain can divide it into subdomains.
10 Each of those subdomains can be delegated to other organizations, in which case the organization receiving the delegation becomes responsible for maintaining all the data in that subdomain. The machines or programs executing on those machines that store information about the domain namespace are often called name servers. Name servers generally have complete information about some part
15 of the domain namespace, namely, a zone. In such case, the name server is then said to have authority for that zone. A name server can also maintain cached information – or non-authoritative information, which is obtained by communicating with an authoritative name server. A zone contains the domain names and data that a domain contains, except for domain names and data that are
20 delegated elsewhere. If a subdomain of the domain has not been delegated, the zone also contains the domain names and data in the subdomain.

For purposes of the present invention, it is assumed that the enterprise operates (or has operated) DNS servers, that those servers manage data held in zones, and that each zone is a complete database for a particular "pruned" subtree
25 of the domain space. As noted above, the zone data is authoritative, in that it contains an accurate (typically, the most accurate) representation available for a particular set of hostnames. At least one host in the enterprise stores enterprise content that the enterprise desires to make available to permitted end users (enterprise or partner employees). The enterprises use the existing DNS

infrastructure to resolve client requests for content at that host. Figure 3 illustrates the known prior art. In this example, it is assumed that the user desires to retrieve the enterprise internal homepage, which is a page located, in this example, at <http://hr.company.com/homepage.html>. To this end, the user, or client 300, interacts with DNS via a resolver 302, typically residing on the same machine. As is well-known, a resolver is a client process that accesses a name server. Programs running on a host that need information from a domain name space use the resolver. In DNS BIND, for example, the resolver is a set of library routines that query a name server, interpret responses, and return the information to the program/process that requested it. Thus, with reference to Figure 3, the resolver 302 interacts with the enterprise name server 304 to resolve client queries, which are then returned to the client. In the example shown in Figure 3, the client resolver 302 makes a DNS request (in this example "who is hr.company.com?") to the corporate DNS name server 304, which returns a reply with an actual IP address y.y.y.y of the host on which the document is stored. The client resolver 302 then uses a DNS A record to build a connection to the actual web server 306, and issues an appropriate HTTP request to the web server 306 (e.g., via a GET <http://hr.company.com/homepage.html>). The web server 306 responds with the content reply to complete the request.

According to the invention, the existing DNS infrastructure is augmented with one or more ECDN zone servers, and requests for content to be delivered over the ECDN are resolved through the one or more zone servers. In the simplest embodiment, this function is achieved by identifying a given enterprise domain to be published to the ECDN, and having a network administrator (or a person acting on the administrator's behalf) modify an existing DNS name record for that enterprise domain. In this embodiment, the administrator CNAMEs the enterprise domain to a domain identified by a so-called "virtual" zone that is managed on an ECDN zone server (and, thus, forms part of the ECDN). In particular, this process enables the enterprise domain to be managed (by the

enterprise or a third party acting on the enterprise's behalf) as part of the ECDN, typically by the ECDN zone server on one of the intelligent nodes (or some other machine, such as a client name server). A similar process is repeated for every new domain or subdomain being distributed by the ECDN. Any client
5 downloading a URL referring to the enterprise domain then gets directed to the ECDN zone server, which directs the client to an ECDN IN. The client passes a host header to the ECDN IN to identify the actual host.

Figure 4 illustrates an embodiment of the present invention. DNS requests typically occur over UDP, while HTTP requests typically occur over TCP. The
10 enterprise is identified as "akamai" for illustrative purposes only. In this example, the client resolver 402 makes a DNS request for "hr.akamai.com" (step (1)) to the configured name server 404. This name server may be part of the legacy DNS infrastructure, as described above. This same name server 404 (running a suitable DNS BIND version) looks up hr.akamai.com and sees that the hostname resides
15 within the zone configured to be within the ECDN (potentially after first resolving a CNAME from hr.akamai.com to itself for hr.ecdn.akamai.com). At step (3), the query is then referred to an ECDN name server 406, with the ECDN name server being configured to be authoritative for the ecdn.akamai.com zone. Some of the steps are compressed in the figure for clarity. At step (4), the ECDN zone server
20 406 could do one of three things: reply with an IP address x.x.x.x, pass the query to a different zone for hr.ecdn.akamai.com, or CNAME the request to a different hostname all together (restarting this process). In the drawing, the ECDN zone server has returned the IP address. Thus, at step (5), the client makes an HTTP request to the ECDN IN 408 at IP address x.x.x.x with a host header: identifying
25 the host hr.akamai.com. At steps (6), (7) and (8), the ECDN content server handles the HTTP portion of the request in the standard fashion for a surrogate. The content is served from the origin server 406.

The technique shown in Figure 4 provides significant advantages. According to the invention, when a DNS query arrives for a hostname known to

be associated with given content within the control of (i.e., managed by) the ECDN, the DNS server that receives the query preferably responds in one of three (3) ways: (a) handing back an IP address, e.g., for an ECDN intelligent node that knows how to obtain the requested content from a surrogate or origin server (as illustrated in Figure 4); (b) executing a zone referral to a next (lower) level in a zone hierarchy (as defined in standard DNS), or (c) CNAMing to another hostname, thereby essentially restarting the lookup procedure. This new CNAME cause the resolution process to start back at the root and resolve a new path, probably along a different path in the hierarchy. At any particular level in the zone hierarchy, there is preferably an associated zone server. Zone servers may operate on the same machines but typically have different IP addresses. Each zone server preferably executes logic that applies the requested hostname against a map which, using known techniques, may be generated from given (static, dynamic, internally-generated or third party-sourced) performance metrics. Thus, a given name query to ECDN-managed content may be serviced in coordination with various sources of distributed network intelligence. As a result, the invention provides for a distributed, dynamic globally load balanced name service for an enterprise CDN.

From a logical view, hierarchical DNS zones for resolving a particular query within a content delivery network are represented according to the present invention as illustrated in Figure 5. In this example, it is assumed that an Internet CDN service provider operates over a namespace that is referred to as "akamai." This namespace includes a subdomain of *.ecdn.akamai.com, which, as noted above, is the virtual zone. In this example, the CDNSP manages the zone on the enterprise's behalf. In a normal name query, a request for foo.akamai.com returns some IP address, which is normally fixed. With web hosting, as is well-known, virtual hostnames are often used to allow a single address to handle multiple hostnames. For load balancing, a round robin technique is often used, allowing for one name to be spread across multiple addresses. For virtual zones, preferably there is a many-to-many mapping. In particular, one hostname is spread across a

mesh of servers, and one server may handle any number of hostnames. Thus, for example, in a simple case, one could create three (3) hostnames (x, y, and z) and have four (4) servers (having IP addresses: 1.1.1.1, 1.1.1.2, 1.1.1.3, 1.1.1.4). Any hostname query could resolve to any server address, giving twelve (12) combinations of responses to any particular query. Thus, when the server for this virtual zone gets a request, it replies with any of the 12 possible addresses, although significant performance advantages could be gained if it were to use a more intelligent approach. In particular, it is also well-known that geographical and topological location information can be used to send clients to the closest address, and server load can be accounted for as well to try and optimize client response times. A combination of these (with other factors) can be used to determine the best server for handling a particular request.

As described above, so-called Canonical Name (CNAME) records define an alias for an official hostname and are part of the DNS standard. In general, CNAMEing may be thought of as a means for DNS redirection of a client query. As illustrated above, DNS virtual zones and CNAMEs are used with standard surrogate servers to manage ECDN content in a given enterprise.

The CNAME technique may be implemented any number of arbitrary times, while zone referral preferably occurs hierarchically. Generalizing, according to the invention, a given domain may be of the form “*.zone2.zone1.enterprise.com” where “zone1” represents a first or “top” level ECDN zone, “zone2” represents a second or “lower” level ECDN subzone of the top level zone, and so forth. As indicated in Figure 6, a first level ECDN zone server 600 would handle top level resolutions, and one or more second level ECDN zone servers 602 would handle the next level resolutions, and so forth. The wildcard “*” in the above example indicates that further zone levels may be provided as well, with each level of the zone hierarchy then inheriting the namespace of the levels above it. Thus, the “virtual” nature of the invention is due to the fact that at no time is there necessarily a fixed set of possible hostnames –

they may be, and most likely are, produced on-the-fly. Thus, for example, if the enterprise uses an Oracle database for storing financial data, the network administrator may generate domains such as "financials.ecdn.company.com," while the ECDN itself may be creating objects like

- 5 "component.oracle.financials.ecdn.company.com," or
"a.b.c.oracle.financials.ecdn.company.com" etc. Continuing with this example, if a particular ECDN zone server is provided with a name it does not fully recognize, it resolves that portion that it can recognize. Thus, while a first-level ECDN DNS server might do a rough cut of performance and refer a request to another zone, the
10 next zone may have a finer grain metric used for next step resolution.

- According to the invention, and as illustrated in Figure 6, assume that given enterprise content has been published to the ECDN and, in particular, to be managed by a second level zone server "hr.ecdn.akamai.com." In this example, a request for resolution of hr.akamai.com is first directed (through a CNAME) to the
15 ecdn.akamai.com top level zone server 600, which then redirects the request (again through a CNAME) to one of the hr.ecdn.akamai.com second level zone servers 602. The second level name servers may be load balanced, e.g., based on network information collected at the zone server pair about connectivity between the remote office and possible locations of the backend servers that host the
20 desired content, as well state information about the load on those servers. The second level name server returns the IP address of an optimal backend server based on the available metrics (e.g., server load, network connectivity, etc.). The use of hierarchical zones in this manner enables the enterprise to build and maintain more relevant "local" information about how client requests should be
25 mapped, thus obviating the building, maintenance and delivery of enterprise-wide request routing maps.

According to another feature of this invention, a software agent is executed in or in conjunction with a DNS server to create "zones" dynamically. In this approach, the agent, in effect, provides hints to the enterprise DNS name server as

to which enterprise domains would be beneficially cached, as well as the addresses of near-by (local) ECDN intelligent nodes. Figure 7 illustrates the agent-hinted CNAME technique. To create such as agent, the enterprise may keep a lightweight history databases within a DNS server, recording usage patterns of A (address) record lookup requests. Building heuristics around access patterns or looking for high or relatively increased lookups on a per domain basis might provide candidate hostnames. This would identify high usage HTTP sites. The hints might also come from firewalls, third party updates, in-line L4 switches, or from some other entity listening on port 80 (e.g., perhaps via an egress switch SPAN port). When a domain is identified as potentially interesting from a usage perspective, instead of returning the address of the actual domain, the address of a nearby ECDN IN may be substituted dynamically. After traffic dies down, the address could revert back to standard behavior. Thus, if the "hot" domain is "hr.akamai.com," the DNS name server dynamically CNAMEs the domain to the local ECDN zone server "ecdn.akamai.com" as needed. All traffic to the hr.akamai.com would then get sent to the close-by ECDN IN for handling. Agent interaction may be accomplished as a server plugin, through WCCP 2.0 transparent interception of DNS traffic, promiscuous DNS port snooping (SPAN mode), or providing an augmented DNS name server. An interface may be provided to notify the agent that all *.ecdn.enterprise.com servers are part of the ECDN and should be sent towards the appropriate servers.

The present invention provides numerous advantages over the prior art. According to the invention, a domain to be published to the ECDN is associated (either by static configuration or dynamically) with a set of one or more enterprise zones configurable in a hierarchy. When a request for content is received from a given client in the enterprise (or otherwise accessible through the firewall), a name query to the enterprise's legacy DNS is directed (preferably via CNAME) to a set of one or more ECDN zone servers. A given ECDN zone server associates a name query with an ECDN intelligent node (or to a given content server). This

association may be based on given conditions within the network, server load conditions, or some combination of various known metrics. With ECDN virtual zones, maps that associate a given request to a server may be local or global, static or dynamic. If multiple zones are used in a hierarchy, a first level ECDN zone server redirects the name query to a second level ECDN zone server, and so on,
5 until the appropriate response is generated. ECDN zone servers may reside on the same machine or with existing DNS infrastructure or other machines in the enterprise. This enables the present invention to be readily integrated into an existing or legacy DNS solution with minimal reconfiguration.

10 As described and illustrated above, a given zone may have a sub-zone and, according to the invention, that sub-zone may be managed by the enterprise or even by a third party such as an Internet content delivery network (ICDN) service provider. As an example of the latter scenario, assume that the enterprise outsources its streaming video delivery to the ICDN service provider. To this end,
15 a host lookup for, say, allhands.streamingvideo.ecdn.company.com may be configured either statically or dynamically to a zone managed by a CDN overlay network. This has the effect of transparently allowing delivery to be seamlessly tied internally and externally via the same mechanism.

Figure 8 illustrates another example of how an enterprise may hand off
20 management of a virtual zone to another entity, in this case an enterprise business partner. In this example, producer.com is the enterprise and consumer.com is the enterprise business partner having an end user. Producer.com desires to make given content available to the consumer.com end user, however, the producer.com namespace on which the content is published is not resolvable through public
25 DNS and the document is not available publicly. Both producer.com and consumer.com are assumed to have implemented an ECDN system behind each of their respective firewalls, and each entity has a private namespace that they respectively manage. The goal is to make producer.com's namespace selectively available to an end user in consumer.com without, at the same time, exposing

producer.com's namespace to the public DNS. Each ECDN deploys a relay machine which, preferably, is a combination of a surrogate origin server and an ECDN zone server. Separate machines/processes may be used, or the processes may be integrated or included with other functionality. Preferably, each relay

5 machine is located within the enterprise DMZ, as illustrated in the drawing. The relay machine inside consumer.com does not know how to resolve a virtual zone associated with producer.com, but that machine does know that the relay machine inside producer.com's firewall does know how to make such a resolution.

Likewise, the relay machine inside producer.com does not know how to resolve a

10 virtual zone associated with consumer.com, but that machine does know that relay machine inside consumer.com's firewall make know how to make such a resolution. As is well-known, boxes within a DMZ can be selectively allowed to communicate with one another while filtering other potentially hostile communications.

15 Referring back to Figure 8, assume that the document in question is referenced within consumer.com by a836.asp.producer.com/object.html. In this example, the zone "asp.producer.com" has been designated within consumer.com as being owned or managed by relay.consumer.com. Thus, relay.consumer.com will act as a surrogate for that document. When the consumer.com relay gets the

20 host header a836.asp.producer.com, the relay machine uses that information to connect to the relay machine in producer.com. The relay machine in producer.com gets the HTTP request and it is configured to rewrite any host header (associated with a836.asp.producer.com in this example) to foo.producer.com. The relay machine at relay.producer.com then acts as a

25 surrogate origin server for the web server at foo.producer.com, which hosts the document. The web server then returns the document via the surrogate origin server relay.producer.com, and that server rewrites the response header to point back to itself. The relay machine in producer.com then returns the data to the relay machine in consumer.com, which performs similar surrogate processing and

returns the document to the requesting end user. This completes the processing. Thus, one of ordinary skill in the art will appreciate that, in this manner, producer.com has handed off management of the relay.producer.com virtual zone to consumer.com. This arrangement creates a VPN-like connection between the
5 two namespaces to facilitate the document sharing. As illustrated, preferably the communications between the two relay machines are made over secure link using known techniques.

Having thus described my invention, the following sets forth what I now claim.

10

CLAIMS

1. A content delivery method operative in an enterprise, comprising:
establishing an enterprise content delivery network within the enterprise,
wherein the enterprise content delivery network has associated therewith a set of
5 one or more zones organized into a hierarchy; and
in response to a name query for a hostname known to be associated with
the given content being managed by the enterprise content delivery network,
returning a response to the name query;
wherein the response selectively (a) identifies a given node in the
10 enterprise content delivery network that can obtain the given content from a server
within the enterprise, (b) causes a referral to a zone within the hierarchy, or (c)
relays the name query to another path.
2. The content delivery method as described in Claim 1 wherein, if
15 the response identifies the given node, having the given node obtain the given
content from the server.
3. The content delivery method as described in Claim 1 wherein, if
the response causes a zone referral, attempting to resolve the name query at a next
20 level within the hierarchy.
4. The content delivery method as described in Claim 1 wherein, if
the response is relayed to another path, attempting to resolve the name query along
that path.
25
5. The content delivery method as described in Claim 1 further
including the step
publishing the given content into the enterprise content delivery network.

6. The content delivery method as described in Claim 5 wherein the step of publishing comprises:

identifying the hostname to be published to the enterprise content delivery network;

5 aliasing the hostname to a given zone in the hierarchy.

7. The content delivery method as described in Claim 6 wherein the aliasing step CNAMEs the hostname to a domain in which the virtual zone is prepended to at least a portion of the hostname.

10

8. The content delivery method as described in Claim 7 wherein the CNAME is generated statically.

9. The content delivery method as described in Claim 7 wherein the CNAME is generated dynamically.

15

10. The content delivery method as described in Claim 1 wherein the name query is resolved to the IP address of the given node as a function of given network or server metrics.

20

11. The content delivery method as described in Claim 10 wherein the IP address of the given node is generated dynamically.

12. The content delivery method as described in Claim 1 wherein the enterprise content delivery network is part of an Internet content delivery network (ICDN).

25

13. The content delivery method as described in Claim 11 wherein the enterprise content delivery network is managed as part of the ICDN.

1/5




Figure 1

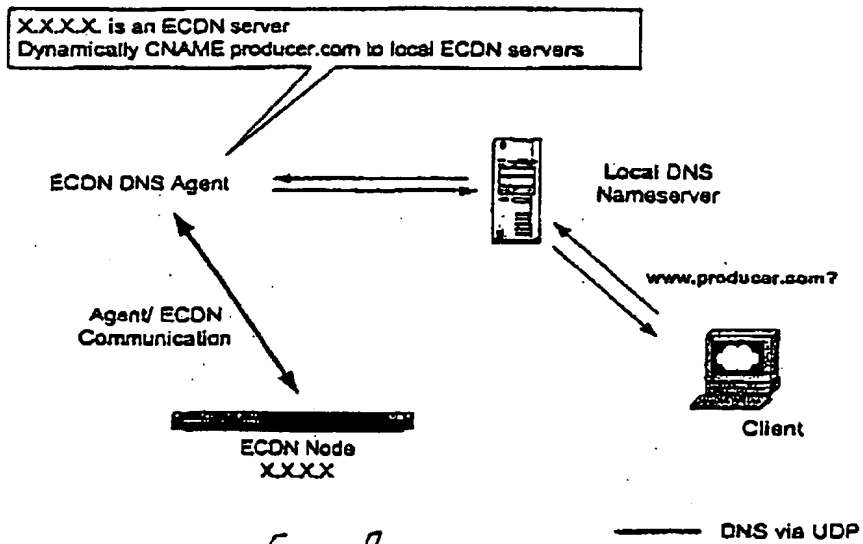


Figure 7

2/5

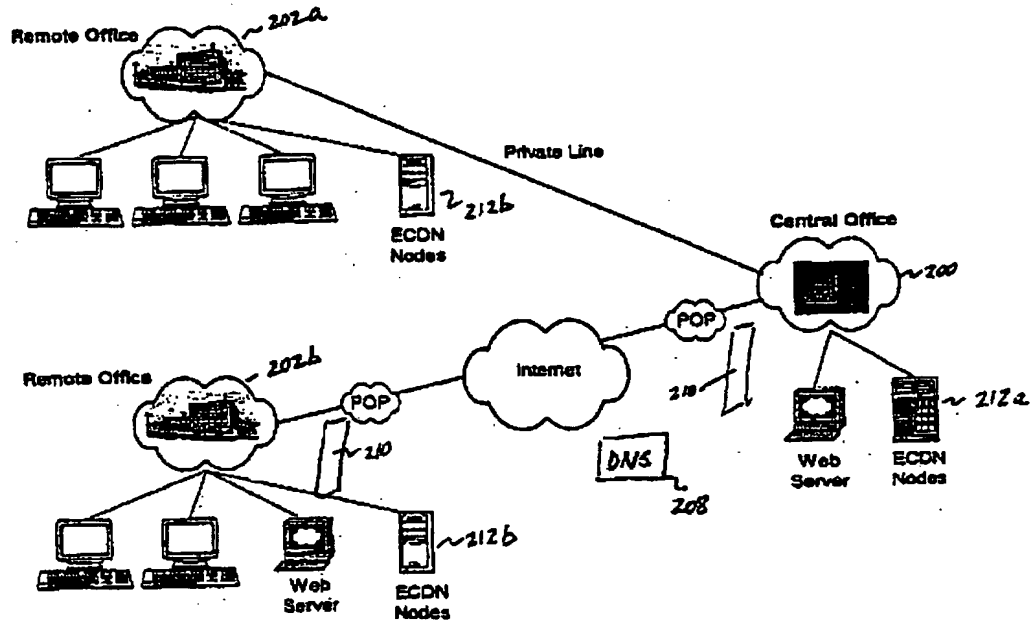


Figure 2

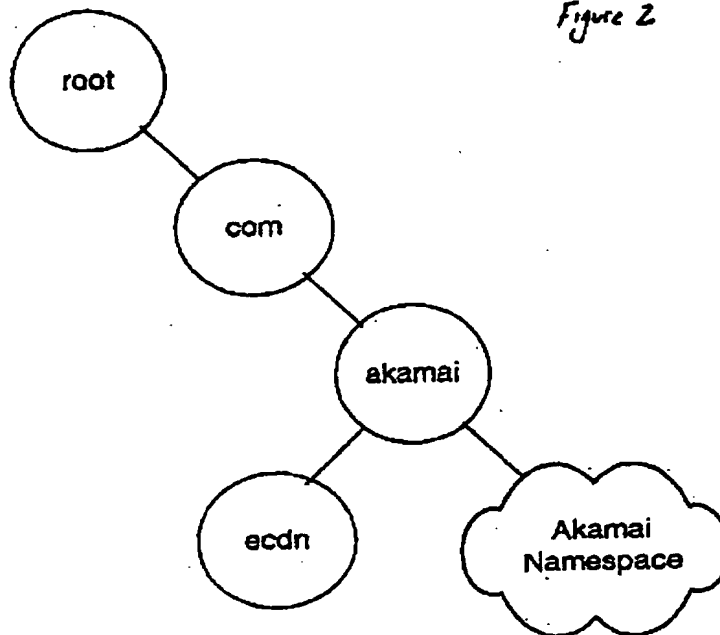



Figure 5


3/5



- 1) Client resolver makes DNS request to name server (NS)
- 2) It gets a reply with the actual address (y.y.y.y).
- 3) Client uses A record, to build connection to web server, and issues HTTP request
- 4) Web server replies with content

Figure 3

4/5



- 1) Client resolver makes DNS request to its name server (NS)
- 2) Local NS looks up `hr.akamai.com`, and so it returns a CNAME redirect to `ecdn.akamai.com` (the shorter name was simply a matter of convenience)
- 3) Client resolver makes DNS request to ecdn zone NS, for hostname `hr.ecdn`
- 4) ECDN zone server replies with an address `x.x.x.x`, or could optionally pass it to another ECDN zone server for `hr` or could CNAME to another hostname
- 5) Assuming it resolves, the client makes an http request to `x.x.x.x` for with Host: `hr.akamai.com`
- 6,7,8,9) ECDN is configured to be a surrogate for `origin-hr.akamai.com`, and satisfies the HTTP request accordingly

Figure 4

5/5

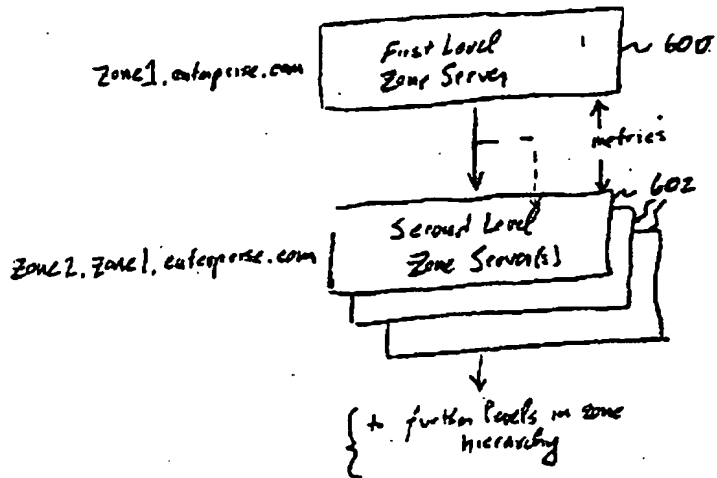


Figure 6


Partner Access to
Enterprise ECDN

Figure 8

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
6 September 2002 (06.09.2002)

PCT

(10) International Publication Number
WO 02/069608 A3

(51) International Patent Classification⁷: **H04L 29/12**,
29/06

(21) International Application Number: **PCT/US02/01079**

(22) International Filing Date: 16 January 2002 (16.01.2002)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/262,171 16 January 2001 (16.01.2001) US

(71) Applicant (for all designated States except US): **AKA-
MAI TECHNOLOGIES, INC.** [US/US]; 500 Technology
Square, Cambridge, MA 02139 (US).

(72) Inventor; and

(75) Inventor/Applicant (for US only): **NEERDAELS**,
Charles, Joseph [US/US]; 622 Bayview Drive, Aptos, CA
95003 (US).

(74) Agent: **JUDSON, David, H.**; Akamai Technologies, Inc.,
500 Technology Square, Cambridge, MA 02139 (US).

(81) Designated States (national): AE, AG, AL, AM, AT, AU,
AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU,

CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH,
GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC,
LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW,
MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG,
SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ,
VN, YU, ZA, ZM, ZW.

(84) Designated States (regional): ARIPO patent (GH, GM,
KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW),
Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),
European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR,
GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent
(BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR,
NE, SN, TD, TG).

Declaration under Rule 4.17:

— of inventorship (Rule 4.17(iv)) for US only

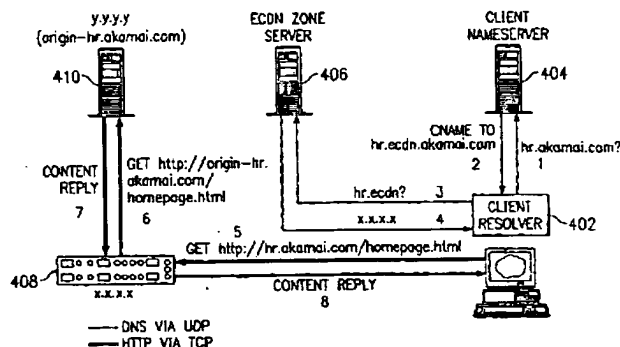
Published:

— with international search report

(88) Date of publication of the international search report:
9 January 2003

For two-letter codes and other abbreviations, refer to the "Guid-
ance Notes on Codes and Abbreviations" appearing at the begin-
ning of each regular issue of the PCT Gazette.

(54) Title: USING VIRTUAL DOMAIN NAME SERVICE (DNS) ZONES FOR ENTERPRISE CONTENT DELIVERY



(57) Abstract: According to the invention, a domain to be published to an enterprise ECDN is associated (either by static configuration or dynamically) with a set of one or more enterprise zones configurable in a hierarchy. When a DNS query arrives for a hostname known to be associated with given content within the control of the ECDN, a DNS server preferably responds in one of three (3) ways; handing back an IP address, e.g., for an ECDN intelligent node that knows how to obtain the requested content from a surrogate or origin server; executing a zone referral to a next (lower) level name server in a zone hierarchy, or CNAMing to another hostname, thereby essentially restarting the lookup procedure. In the latter case, this new CNAME causes the resolution process to start back at the root and resolve a new path, probably along a different path in the hierarchy.

WO 02/069608 A3

INTERNATIONAL SEARCH REPORT

Int'l Application No
PCT/US 02/01079

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04L29/12 H04L29/06

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 99 40514 A (BALTER JAMES A ;FARBER DAVID A (US); GREER RICHARD E (US); SWART A) 12 August 1999 (1999-08-12) abstract page 2, line 18 -page 3, line 17 page 4, line 1 -page 5, line 24 page 6, line 24 -page 8, line 12 page 12, line 9 -page 12, line 15 page 13, line 3 -page 13, line 13	1-4, 10-13
X	US 6 108 703 A (LEWIN DANIEL M ET AL) 22 August 2000 (2000-08-22) abstract column 3, line 24 -column 4, line 22 -/-	1-4, 10-13

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents :

A document defining the general state of the art which is not considered to be of particular relevance

E earlier document but published on or after the international filing date

L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

O document referring to an oral disclosure, use, exhibition or other means

P document published prior to the international filing date but later than the priority date claimed

T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

G document member of the same patent family

Date of the actual completion of the international search

5 September 2002

Date of mailing of the international search report

13/09/2002

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Goller, W

INTERNATIONAL SEARCH REPORT

In national Application No
PCT/US 02/01079

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT		
Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 0 817 444 A (SUN MICROSYSTEMS INC) 7 January 1998 (1998-01-07) abstract column 2, line 20 -column 3, line 41 column 5, line 50 -column 6, line 57 column 7, line 30 -column 7, line 40	1,2, 10-13
X	WO 00 74347 A (ENTERA INC) 7 December 2000 (2000-12-07) abstract page 5, line 4 -page 6, line 18 page 9, line 2 -page 9, line 24	1-4, 10-13

Form PCT/ISA/210 (continuation of second sheet) (July 1992)

INTERNATIONAL SEARCH REPORT

Information on patent family members

Int. Application No

PCT/US 02/01079

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
WO 9940514	A	12-08-1999	US 6185598 B1	06-02-2001
			AU 2652999 A	23-08-1999
			CA 2320261 A1	12-08-1999
			EP 1143337 A1	10-10-2001
			EP 1053524 A1	22-11-2000
			JP 2002503001 T	29-01-2002
			NO 20004010 A	10-10-2000
			US 2002049857 A1	25-04-2002
			US 2002099850 A1	25-07-2002
			WO 9940514 A1	12-08-1999
			US 2001056500 A1	27-12-2001
US 6108703	A	22-08-2000	AU 4995299 A	07-02-2000
			BR 9912001 A	04-12-2001
			CN 1312923 T	12-09-2001
			DE 1125219 T1	23-05-2002
			EP 1125219 A1	22-08-2001
			JP 2002520735 T	09-07-2002
			WO 0004458 A1	27-01-2000
			US 2002099850 A1	25-07-2002
EP 0817444	A	07-01-1998	US 6154777 A	28-11-2000
			EP 0817444 A2	07-01-1998
			JP 10126445 A	15-05-1998
WO 0074347	A	07-12-2000	AU 4661600 A	18-12-2000
			WO 0074347 A1	07-12-2000

Form PCT/ISA/210 (patent family annex) (July 1992)

CORRECTED VERSION

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
6 September 2002 (06.09.2002)

PCT

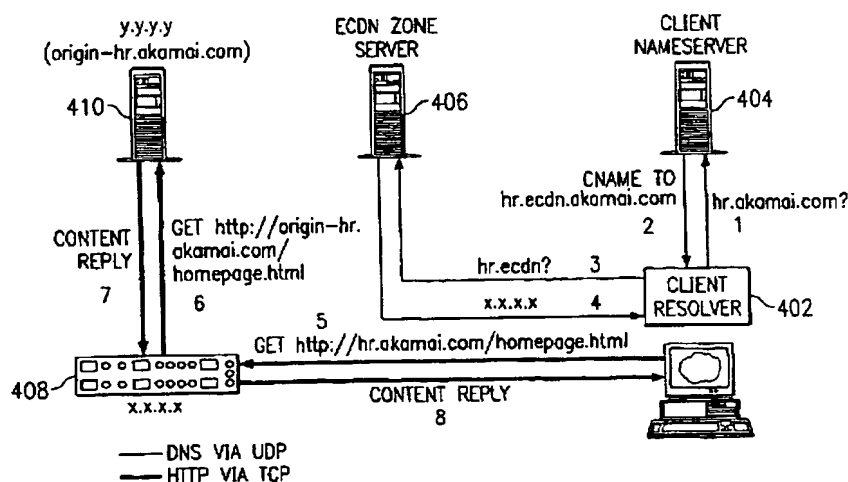
(10) International Publication Number
WO 02/069608 A3

- (51) International Patent Classification?: **H04L 29/12**, 29/06 (74) Agent: JUDSON, David, H.; Akamai Technologies, Inc., 500 Technology Square, Cambridge, MA 02139 (US).
- (21) International Application Number: PCT/US02/01079 (81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZM, ZW.
- (22) International Filing Date: 16 January 2002 (16.01.2002)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data: 60/262,171 16 January 2001 (16.01.2001) US (84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).
- (71) Applicant (for all designated States except US): **AKA-MAI TECHNOLOGIES, INC.** [US/US]; 500 Technology Square, Cambridge, MA 02139 (US).
- (72) Inventor; and
- (75) Inventor/Applicant (for US only): **NEERDAELS, Charles, Joseph** [US/US]; 622 Bayview Drive, Aptos, CA 95003 (US).

Declaration under Rule 4.17:
— of inventorship (Rule 4.17(iv)) for US only

[Continued on next page]

(54) Title: USING VIRTUAL DOMAIN NAME SERVICE (DNS) ZONES FOR ENTERPRISE CONTENT DELIVERY



(57) Abstract: According to the invention, a domain to be published to an enterprise ECDN is associated (either by static configuration or dynamically) with a set of one or more enterprise zones configurable in a hierarchy. When a DNS query arrives for a hostname known to be associated with given content within the control of the ECDN, a DNS server preferably responds in one of three (3) ways; handing back an IP address, e.g., for an ECDN intelligent node that knows how to obtain the requested content from a surrogate or origin server; executing a zone referral to a next (lower) level name server in a zone hierarchy, or CNAMing to another hostname, thereby essentially restarting the lookup procedure. In the latter case, this new CNAME causes the resolution process to start back at the root and resolve a new path, probably along a different path in the hierarchy.

WO 02/069608 A3



Published:

— *with international search report*

(15) Information about Correction:

see PCT Gazette No. 10/2003 of 6 March 2003, Section II

(88) Date of publication of the international search report:

9 January 2003

(48) Date of publication of this corrected version:

6 March 2003

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

USING VIRTUAL DOMAIN NAME SERVICE (DNS) ZONES FOR ENTERPRISE CONTENT DELIVERY

BACKGROUND OF THE INVENTION

This application is based on and claims priority from Provisional

5 Application Serial No. 60/262,171 filed January 16, 2001.

Technical Field

The present invention relates generally to content delivery and management within a private enterprise network.

Description of the Related Art

10 It is well-known to deliver digital content (e.g., HTTP content, streaming media and applications) using an Internet content delivery network (ICDN). A content delivery network or "CDN" is a network of geographically distributed content delivery nodes that are arranged for efficient delivery of content on behalf of third party content providers. A request from a requesting end user for given
15 content is directed to a "best" replica, where "best" usually means that the item is served to the client quickly compared to the time it would take to fetch it from the content provider origin server.

Typically, a CDN is implemented as a combination of a content delivery infrastructure, a request-routing mechanism, and a distribution infrastructure. The
20 content delivery infrastructure usually comprises a set of "surrogate" origin servers that are located at strategic locations (e.g., Internet network access points, Internet Points of Presence, and the like) for delivering copies of content to requesting end users. The request-routing mechanism allocates servers in the content delivery infrastructure to requesting clients in a way that, for web content delivery,
25 minimizes a given client's response time and, for streaming media delivery, provides for the highest quality. The distribution infrastructure consists of on-demand or push-based mechanisms that move content from the origin server to the surrogates. An effective CDN serves frequently-accessed content from a surrogate that is optimal for a given requesting client. In a typical CDN, a single service

provider operates the request-routers, the surrogates, and the content distributors. In addition, that service provider establishes business relationships with content publishers and acts on behalf of their origin server sites to provide a distributed delivery system. A well-known commercial CDN service that provides web
5 content and media streaming is provided by Akamai Technologies, Inc. of Cambridge, Massachusetts.

Enterprises have begun to explore the desirability of implementing content delivery infrastructures to address several problems. Currently, enterprise users typically experience slow and expensive access to Internet content. Slow access to
10 business critical data available on the Internet hurts productivity, and the cost of providing good access, e.g., by building bigger networks and by deploying and managing caching infrastructure, is large. In addition, many IT organizations cannot deliver the required quality of service for Internet content delivery due to lack of talent and expertise. Yet another reason corporations are exploring CDNs
15 is because of the slow, expensive and often cumbersome access to and within the entity's intranet. As corporate intranets quickly become a critical component of business process in many large companies, fast and efficient access to the data and applications on the intranet is a high priority for many IT departments. Nevertheless, current intranet delivery solutions are inadequate, and solving the
20 problems, e.g., by building bigger internal networks, deploying and managing caches, and distributing application front ends, is extremely expensive. To address these deficiencies, several large software vendors are attempting to build ecosystems to provide web-based front ends to many enterprise applications, however, distributing these application from front ends efficiently, in of itself, will
25 be a critical IT problem that current technologies do not addresss. Finally, enterprises are considering CDN technology due to slow, expensive access to business partner applications and information provided by current techniques and solutions. Business-to-business applications (such as ordering, inventory, and pricing management) between business partners is done today by linking partners

with a physical network. These applications are moving to the Internet/intranet, and the need to link business partners together in an efficient way with web-based front ends is another critical IT problem that is not addressed by today's solutions.

BRIEF SUMMARY OF THE INVENTION

5 It is a general object of the invention to define and implement one or more virtual zones within an enterprise namespace to facilitate enterprise content delivery behind a corporate firewall.

 It is another general object of the present invention to provide an enterprise content delivery network that coexists with existing or "legacy" Domain Name
10 Service (DNS) infrastructure to facilitate mapping of requests for enterprise resources to surrogate servers, e.g., common caching appliances, application servers, distributed storage and database management systems, and streaming media servers.

 It is yet another general object of the invention to define and implement
15 one or more so-called "virtual" zones within an enterprise namespace to facilitate content delivery behind a corporate firewall over an enterprise content delivery network (ECDN).

 Another object of the present invention to enable an ECDN DNS to coexist with an existing enterprise DNS and to enable content delivery with only minimal
20 configuration changes to the existing infrastructure using virtual zones. A given host or sub-domain to be published to the ECDN is aliased into a virtual zone namespace preferably using a DNS Canonical Name (CNAME).

 A still further object of the invention is to implement an enterprise content delivery network wherein any arbitrary hierarchical namespace can be
25 implemented and wherein each layer of the namespace inherits the namespace above.

 According to the invention, a domain to be published to an enterprise ECDN is associated (either by static configuration or dynamically) with a set of one or more enterprise zones configurable in a hierarchy. When a DNS query

arrives for a hostname known to be associated with given content within the control of the ECDN, a DNS server preferably responds in one of three (3) ways: (a) handing back an IP address, e.g., for an ECDN intelligent node that knows how to obtain the requested content from a surrogate or origin server; (b) executing a zone referral to a next (lower) level name server in a zone hierarchy, or (c) CNAMing to another hostname, thereby essentially restarting the lookup procedure. In the latter case, this new CNAME causes the resolution process to start back at the root and resolve a new path, probably along a different path in the hierarchy. At any particular level in the zone hierarchy, preferably there is an associated zone server. That server preferably executes logic that applies the requested hostname against a map, which, using known techniques, may be generated from given (static, dynamic, internally-generated or third party-sourced) performance metrics. Thus, a given name query to ECDN-managed content may be serviced in coordination with various sources of distributed network intelligence. As a result, the invention provides for a distributed, dynamic globally load balanced name service.

According to a more specific aspect of the invention, a domain to be published to the ECDN is associated (either by static configuration or dynamically) with a set of one or more enterprise zones configurable in a hierarchy. When a request for content is received from a given client in the enterprise (or otherwise accessible through the firewall), a name query to the enterprise's legacy DNS is directed to the configured ECDN zone primary server. This primary ECDN zone server associates, and relays via CNAME, a name query with an ECDN intelligent node (or resolves directly to a given content origin server). This association may be based on given conditions within the network, server load conditions, or some combination of various known metrics. With ECDN virtual zones, maps that associate a given request to a server may be local or global, static or dynamic. While the maps typically are generated within the ECDN infrastructure, they could also be affected by external agents providing

hints or modifications. If multiple zones are used in a hierarchy, a first level ECDN zone server directs the name query to a second level ECDN zone server, and so on, until the appropriate response is generated. At any one level the particular zone server need not be able to resolve the entire name; rather, the particular zone server need only be able to resolve the portion required for directing the next level of resolution. For example: the zone server for the namespace containing b in the hostname a.b.c.d.e may have no idea where the final resolution will end up, but it should know which set of servers may be authoritative for the namespace containing c. Because DNS readily allows one to “carve out” zones in the namespace of a particular domain (e.g., ecdn.company.com can be managed by a different authoritative DNS server than the rest of company.com), the creation of the first level ECDN zone on the existing DNS infrastructure allows the two systems to coexist and undergo modification in parallel without interference with one another. This enables the present invention to be readily integrated into an existing or legacy DNS solution with minimal reconfiguration.

Thus, according to the invention, an existing DNS infrastructure is augmented with one or more ECDN zone servers, and requests for content to be delivered over the ECDN are resolved through the one or more zone servers. In one embodiment, this function is achieved by identifying a given enterprise domain to be published to the ECDN, and having a network administrator (or a person acting on the administrator's behalf) modify an existing DNS name record for that enterprise domain. In this embodiment, the administrator CNAMEs the enterprise domain to a domain identified by a so-called “virtual zone” that is managed on an ECDN zone server (and, thus, forms part of the ECDN). In particular, this process enables the enterprise domain to be managed (by the enterprise or a third party acting on the enterprise's behalf) as part of the ECDN, typically by the ECDN zone server on one of the intelligent nodes (or some other machine, such as a client name server). A similar process is repeated for every

new domain or subdomain being distributed by the ECDN. Any client downloading a URL referring to the enterprise domain then gets directed to the ECDN zone server, which directs the client to an ECDN intelligent node. The client passes a host header to the ECDN node to identify the actual host.

- 5 The foregoing has outlined some of the more pertinent features of the present invention. These features should be construed to be merely illustrative. Many other beneficial results can be attained by applying the disclosed invention in a different manner or by modifying the invention as will be described.

BRIEF DESCRIPTION OF THE DRAWINGS

- 10 Figure 1 is a block diagram of an enterprise in which the inventive enterprise content delivery network (ECDN) may be implemented;

Figure 2 is another enterprise environment in which the present invention may be implemented;

- 15 Figure 3 illustrates a known technique for handling a request for content in an enterprise using a legacy DNS infrastructure;

Figure 4 illustrates a technique for enterprise content delivery using virtual DNS zones according to the present invention;

Figure 5 illustrates a representative DNS namespace in which an enterprise virtual zone is created according to the present invention;

- 20 Figure 6 illustrates a zone hierarchy in which first and second level ECDN zone servers are used to resolve a name query to an enterprise domain that has published to the ECDN; and

Figure 7 illustrates a technique for dynamic zone creation through an agent-hinted CNAME according to the present invention; and

- 25 Figure 8 illustrates a technique for using virtual zones to enable an enterprise partner to access an enterprise ECDN to retrieve intranet content.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

The present invention is implemented within an "enterprise" or "enterprise environment." Referring to Figure 1, a representative enterprise environment in

which the present invention is implemented comprises at least one primary central office 100 (e.g., with central IT support), which is connected to one or more regional datacenters 102, with each datacenter connected to one or more remote branch offices 104 via private line, or via the public Internet (most likely with a virtual private network or "VPN"). The branch offices 104 may be fully (or partially) meshed amongst themselves. Figure 2 illustrates another enterprise computing environment that includes a central office 200 and a pair of remote offices 202. As shown in Figure 2, remote office 202a is connected to the central office 200 over a private line, which refers to a line not generally routable over the public Internet (e.g., frame relay, satellite link, microwave link, or the like), and remote office 202b is connected to the central office 200 over a virtual private network (VPN) typically over the public Internet, or in other known ways. A given remote office may be an office of an enterprise business partner if an appropriate business relationship between the enterprise and the partner exists.

Enterprise networks are generally based on one one of two network types – point-to-point private/leased line, or VPN over public lines (often fully meshed between offices). From a topological view, actual enterprise networks may not operate in a hierarchical sense for connectivity. From a logical management view, however, the servers remain in a relatively weak hierarchy (as illustrated in either Figure 1 or 2) with more valuable and management intensive systems usually existing at more centralized data centers, and very few of these types of systems being deployed in remote branched offices. This architecture is usually a function of where information and management reside, which drives the cost of management for the overall IT network and IT services. Thus, for example, in the more complex enterprise such as shown in Figure 1, regional datacenters tend to have more critical information than branch offices, and the central office tends to have the most critical information. It is assumed for purposes of illustration that the enterprise hosts given internal content that it desires to have stored, cached and delivered to end users in the enterprise. Such internal enterprise content, or

ECDN content, generally is non publicly-available content (in whatever format) that an enterprise desires to make available to permitted users within the enterprise or within a third party partner of the enterprise. An end user typically operates a computer having an Internet browser. The authoritative content and application
5 servers 106 (in Figure 1) and 206 (in Figure 2) being published into the CDN (each identified here as an "ECDN server") can exist in any location, although from a practical perspective they may often be somewhat centrally-located. Such servers are sometimes referred to as origin servers. The enterprise is also assumed to be IP based and to have a Domain Name Service (DNS) 108 (in Figure 1) and
10 208 (in Figure 2). From a security standpoint, the enterprise network manager roughly divides the world of the network into trusted and un-trusted, which usually corresponds to internal and external entities. Firewalls 110 (in Figure 1) and 210 (in Figure 2) are typically configured to allow most outbound traffic, with severe restrictions on inbound traffic. This allows for access to requested services
15 without providing third parties a foothold for intrusion. More sophisticated systems usually create a security entity called a DMZ, which can be thought of as a set of two firewalls, with certain assets like email, DNS, web servers, etc. sitting between them. Each firewall has a different set of filtering rules, with the innermost generally allowing valid traffic from a host within the DMZ to enter the
20 enterprise. Generally, no externally initiated traffic is allowed to pass both sets of filters.

For purposes of illustration, an enterprise content delivery network (ECDN) is provisioned into the enterprise network topology of Figure 2. The ECDN may be implemented in a standalone manner, or it may be part of a larger
25 Internet CDN, in which case the Internet CDN service provider may manage the ECDN on behalf of the enterprise. Various techniques for extending an ICDN into an enterprise are described in copending application Serial No. xx/xxx,xxx, filed January 7, 2002, titled **EXTENDING AN INTERNET CONTENT DELIVERY NETWORK INTO AN ENTERPRISE**, and in copending

application Serial No. yy/yyy,yyy, filed January 7, 2002, titled **EXTENDING AN INTERNET CONTENT DELIVERY NETWORK INTO AN ENTERPRISE ENVIRONMENT BY LOCATING ICDN CONTENT SERVERS TOPOLOGICALLY NEAR AN ENTERPRISE FIREWALL**, which

5 applications are commonly-owned and which are incorporated herein by reference. Referring back to Figure 2, ECDN functionality is provided by locating one or more intelligent nodes (each referred to as an IN) that execute appropriate management, provisioning, monitoring and routing applications to facilitate the distribution, delivery and management of the private enterprise content delivery
10 network. Thus, for example, an IN 212a may be located in the central office 200 and each of the remote offices may include an IN box 212b as indicated. An IN box 212 is typically a computer, namely, a server, having at least one Pentium-class processor, an operating system (e.g., Linux, Windows NT, Windows 2000, or the like), and some amount of disk storage and system memory. Using the IN
15 boxes, the enterprise (or the CDNSP on its behalf) has the ability to map and load balance users inside the enterprise (perhaps as part of the global ICDN), to fetch content from inside the firewall, to collect log information, to deploy software, to distribute live streams, and to provide other ICDN functionality.

Thus, an illustrative enterprise CDN ("ECDN") is comprised of a variety
20 of origin and surrogate servers and a distributed set of intelligence nodes (IN). An origin server can be any server that is intended to respond to a user request for information and includes, for example, common caching appliances, application servers, distributed storage and database management systems, and streaming media servers. The intelligence nodes tie the entire set of distributed nodes and
25 services into a single, virtual service cloud. The ECDN intelligence nodes may have different configurations. Thus, for example, a first type (e.g., IN 212a in central office 200) may be sized (in terms of processing power and memory) for the largest datacenters, and such machines would be aggregation points for monitoring, logging, management and request routing. A second type of

intelligence node (e.g., IN 212b) may be sized for branch office deployment, and such servers would be sized to optimize local behavior. The nodes may be placed in a hierarchy with an IN of the first type acting as a parent. An intelligence node of the second type would typically include a caching engine.

5 Architecturally, an ECDN such as illustrated above is similar to an Internet content delivery network (ICDN) and comprises: HTTP, streaming and DNS server nodes deployed across a customer network. Generalizing, the ECDN provides the enterprise with the following: a namespace and an associated request routing mechanism (to control how to navigate the CDN), storage, caching and
10 streaming (to control how and where are objects stored and distributed), optionally an application environment (to facilitate application extensibility), management and provisioning (to control how the CDN is deployed and managed), and reporting/auditing (to provide visibility into the behavior). An ECDN such as described above provides an application, streaming and content delivery platform
15 so that enterprise content may be delivered and managed in a decentralized fashion while providing a tightly-integrated solution with a customer's storage, network, application and management solutions.

Enterprises today either outsource their Domain Name Service (DNS) or, as described above, run a DNS infrastructure (e.g., one or more DNS servers). A
20 DNS namespace has certain characteristics, as are well-known. Each unit of data in a DNS distributed database is indexed by a hostname. These names are simply leaves in a tree representation, which is often called a domain namespace. Each node in the tree is labeled with a simple name. The fully qualified hostname of any node in the tree is the sequence of labels on the path from that node to the
25 root. A zone is a subtree of the domain name space. Names at the leaves of the tree generally represent individual hosts, whose names resolve to one or more IP addresses. It has become common in instances of virtual hosting within ISPs to have multiple customers share a single machine, therefore having multiple hostnames resolve to a single machine or server farm. CDN technology to date has

additionally made it common for a single hostname to potentially resolve to any one of thousands of machines, though at any given time the set of possible hostnames is configured and fixed. The term zone usually relates to a set of hosts within a single authoritative server's namespace, whereas a single zone could
5 include multiple sub-domains. Thus, e.g., if foo.com and eng.foo.com are both handled by the same nameserver, they are both in the same zone; but, if foo.com's nameserver refers a lookup to a different authority for resolution, they would be considered separate zones.

An organization administrating a domain can divide it into subdomains.
10 Each of those subdomains can be delegated to other organizations, in which case the organization receiving the delegation becomes responsible for maintaining all the data in that subdomain. The machines or programs executing on those machines that store information about the domain namespace are often called name servers. Name servers generally have complete information about some part
15 of the domain namespace, namely, a zone. In such case, the name server is then said to have authority for that zone. A name server can also maintain cached information – or non-authoritative information, which is obtained by communicating with an authoritative name server. A zone contains the domain names and data that a domain contains, except for domain names and data that are
20 delegated elsewhere. If a subdomain of the domain has not been delegated, the zone also contains the domain names and data in the subdomain.

For purposes of the present invention, it is assumed that the enterprise operates (or has operated) DNS servers, that those servers manage data held in zones, and that each zone is a complete database for a particular "pruned" subtree
25 of the domain space. As noted above, the zone data is authoritative, in that it contains an accurate (typically, the most accurate) representation available for a particular set of hostnames. At least one host in the enterprise stores enterprise content that the enterprise desires to make available to permitted end users (enterprise or partner employees). The enterprises use the existing DNS

infrastructure to resolve client requests for content at that host. Figure 3 illustrates the known prior art. In this example, it is assumed that the user desires to retrieve the enterprise internal homepage, which is a page located, in this example, at <http://hr.company.com/homepage.html>. To this end, the user, or client 300, interacts with DNS via a resolver 302, typically residing on the same machine. As is well-known, a resolver is a client process that accesses a name server. Programs running on a host that need information from a domain name space use the resolver. In DNS BIND, for example, the resolver is a set of library routines that query a name server, interpret responses, and return the information to the program/process that requested it. Thus, with reference to Figure 3, the resolver 302 interacts with the enterprise name server 304 to resolve client queries, which are then returned to the client. In the example shown in Figure 3, the client resolver 302 makes a DNS request (in this example "who is hr.company.com?") to the corporate DNS name server 304, which returns a reply with an actual IP address y.y.y.y of the host on which the document is stored. The client resolver 302 then uses a DNS A record to build a connection to the actual web server 306, and issues an appropriate HTTP request to the web server 306 (e.g., via a GET <http://hr.company.com/homepage.html>). The web server 306 responds with the content reply to complete the request.

According to the invention, the existing DNS infrastructure is augmented with one or more ECDN zone servers, and requests for content to be delivered over the ECDN are resolved through the one or more zone servers. In the simplest embodiment, this function is achieved by identifying a given enterprise domain to be published to the ECDN, and having a network administrator (or a person acting on the administrator's behalf) modify an existing DNS name record for that enterprise domain. In this embodiment, the administrator CNAMEs the enterprise domain to a domain identified by a so-called "virtual" zone that is managed on an ECDN zone server (and, thus, forms part of the ECDN). In particular, this process enables the enterprise domain to be managed (by the

enterprise or a third party acting on the enterprise's behalf) as part of the ECDN, typically by the ECDN zone server on one of the intelligent nodes (or some other machine, such as a client name server). A similar process is repeated for every new domain or subdomain being distributed by the ECDN. Any client
5 downloading a URL referring to the enterprise domain then gets directed to the ECDN zone server, which directs the client to an ECDN IN. The client passes a host header to the ECDN IN to identify the actual host.

Figure 4 illustrates an embodiment of the present invention. DNS requests typically occur over UDP, while HTTP requests typically occur over TCP. The
10 enterprise is identified as "akamai" for illustrative purposes only. In this example, the client resolver 402 makes a DNS request for "hr.akamai.com" (step (1)) to the configured name server 404. This name server may be part of the legacy DNS infrastructure, as described above. This same name server 404 (running a suitable DNS BIND version) looks up hr.akamai.com and sees that the hostname resides
15 within the zone configured to be within the ECDN (potentially after first resolving a CNAME from hr.akamai.com to itself for hr.ecdn.akamai.com). At step (3), the query is then referred to an ECDN name server 406, with the ECDN name server being configured to be authoritative for the ecdn.akamai.com zone. Some of the steps are compressed in the figure for clarity. At step (4), the ECDN zone server
20 406 could do one of three things: reply with an IP address x.x.x.x, pass the query to a different zone for hr.ecdn.akamai.com, or CNAME the request to a different hostname all together (restarting this process). In the drawing, the ECDN zone server has returned the IP address. Thus, at step (5), the client makes an HTTP request to the ECDN IN 408 at IP address x.x.x.x with a host header: identifying
25 the host hr.akamai.com. At steps (6), (7) and (8), the ECDN content server handles the HTTP portion of the request in the standard fashion for a surrogate. The content is served from the origin server 406.

The technique shown in Figure 4 provides significant advantages. According to the invention, when a DNS query arrives for a hostname known to

be associated with given content within the control of (i.e., managed by) the ECDN, the DNS server that receives the query preferably responds in one of three (3) ways: (a) handing back an IP address, e.g., for an ECDN intelligent node that knows how to obtain the requested content from a surrogate or origin server (as illustrated in Figure 4); (b) executing a zone referral to a next (lower) level in a zone hierarchy (as defined in standard DNS), or (c) CNAMing to another hostname, thereby essentially restarting the lookup procedure. This new CNAME cause the resolution process to start back at the root and resolve a new path, probably along a different path in the hierarchy. At any particular level in the zone hierarchy, there is preferably an associated zone server. Zone servers may operate on the same machines but typically have different IP addresses. Each zone server preferably executes logic that applies the requested hostname against a map which, using known techniques, may be generated from given (static, dynamic, internally-generated or third party-sourced) performance metrics. Thus, a given name query to ECDN-managed content may be serviced in coordination with various sources of distributed network intelligence. As a result, the invention provides for a distributed, dynamic globally load balanced name service for an enterprise CDN.

From a logical view, hierarchical DNS zones for resolving a particular query within a content delivery network are represented according to the present invention as illustrated in Figure 5. In this example, it is assumed that an Internet CDN service provider operates over a namespace that is referred to as "akamai." This namespace includes a subdomain of *.ecdن.akamai.com, which, as noted above, is the virtual zone. In this example, the CDNSP manages the zone on the enterprise's behalf. In a normal name query, a request for foo.akamai.com returns some IP address, which is normally fixed. With web hosting, as is well-known, virtual hostnames are often used to allow a single address to handle multiple hostnames. For load balancing, a round robin technique is often used, allowing for one name to be spread across multiple addresses. For virtual zones, preferably there is a many-to-many mapping. In particular, one hostname is spread across a

mesh of servers, and one server may handle any number of hostnames. Thus, for example, in a simple case, one could create three (3) hostnames (x, y, and z) and have four (4) servers (having IP addresses: 1.1.1.1, 1.1.1.2, 1.1.1.3, 1.1.1.4). Any hostname query could resolve to any server address, giving twelve (12) combinations of responses to any particular query. Thus, when the server for this virtual zone gets a request, it replies with any of the 12 possible addresses, although significant performance advantages could be gained if it were to use a more intelligent approach. In particular, it is also well-known that geographical and topological location information can be used to send clients to the closest address, and server load can be accounted for as well to try and optimize client response times. A combination of these (with other factors) can be used to determine the best server for handling a particular request.

As described above, so-called Canonical Name (CNAME) records define an alias for an official hostname and are part of the DNS standard. In general, CNAMEing may be thought of as a means for DNS redirection of a client query. As illustrated above, DNS virtual zones and CNAMEs are used with standard surrogate servers to manage ECDN content in a given enterprise.

The CNAME technique may be implemented any number of arbitrary times, while zone referral preferably occurs hierarchically. Generalizing, according to the invention, a given domain may be of the form “*.zone2.zone1.enterprise.com” where “zone1” represents a first or “top” level ECDN zone, “zone2” represents a second or “lower” level ECDN subzone of the top level zone, and so forth. As indicated in Figure 6, a first level ECDN zone server 600 would handle top level resolutions, and one or more second level ECDN zone servers 602 would handle the next level resolutions, and so forth. The wildcard “*” in the above example indicates that further zone levels may be provided as well, with each level of the zone hierarchy then inheriting the namespace of the levels above it. Thus, the “virtual” nature of the invention is due to the fact that at no time is there necessarily a fixed set of possible hostnames –

they may be, and most likely are, produced on-the-fly. Thus, for example, if the enterprise uses an Oracle database for storing financial data, the network administrator may generate domains such as "financials.ecdn.company.com," while the ECDN itself may be creating objects like

5 "component.oracle.financials.ecdn.company.com," or
"a.b.c.oracle.financials.ecdn.company.com" etc. Continuing with this example, if a particular ECDN zone server is provided with a name it does not fully recognize, it resolves that portion that it can recognize. Thus, while a first-level ECDN DNS server might do a rough cut of performance and refer a request to another zone, the
10 next zone may have a finer grain metric used for next step resolution.

According to the invention, and as illustrated in Figure 6, assume that given enterprise content has been published to the ECDN and, in particular, to be managed by a second level zone server "hr.ecdn.akamai.com." In this example, a request for resolution of hr.akamai.com is first directed (through a CNAME) to the
15 ecdn.akamai.com top level zone server 600, which then redirects the request (again through a CNAME) to one of the hr.ecdn.akamai.com second level zone servers 602. The second level name servers may be load balanced, e.g., based on network information collected at the zone server pair about connectivity between the remote office and possible locations of the backend servers that host the
20 desired content, as well state information about the load on those servers. The second level name server returns the IP address of an optimal backend server based on the available metrics (e.g., server load, network connectivity, etc.). The use of hierarchical zones in this manner enables the enterprise to build and maintain more relevant "local" information about how client requests should be
25 mapped, thus obviating the building, maintenance and delivery of enterprise-wide request routing maps.

According to another feature of this invention, a software agent is executed in or in conjunction with a DNS server to create "zones" dynamically. In this approach, the agent, in effect, provides hints to the enterprise DNS name server as

to which enterprise domains would be beneficially cached, as well as the addresses of near-by (local) ECDN intelligent nodes. Figure 7 illustrates the agent-hinted CNAME technique. To create such as agent, the enterprise may keep a lightweight history databases within a DNS server, recording usage patterns of A
5 (address) record lookup requests. Building heuristics around access patterns or looking for high or relatively increased lookups on a per domain basis might provide candidate hostnames. This would identify high usage HTTP sites. The hints might also come from firewalls, third party updates, in-line L4 switches, or from some other entity listening on port 80 (e.g., perhaps via an egress switch
10 SPAN port). When a domain is identified as potentially interesting from a usage perspective, instead of returning the address of the actual domain, the address of a nearby ECDN IN may be substituted dynamically. After traffic dies down, the address could revert back to standard behavior. Thus, if the "hot" domain is "hr.akamai.com," the DNS name server dynamically CNAMEs the domain to the
15 local ECDN zone server "ecdn.akamai.com" as needed. All traffic to the hr.akamai.com would then get sent to the close-by ECDN IN for handling. Agent interaction may be accomplished as a server plugin, through WCCP 2.0 transparent interception of DNS traffic, promiscuous DNS port snooping (SPAN mode), or providing an augmented DNS name server. An interface may be
20 provided to notify the agent that all *.ecdn.enterprise.com servers are part of the ECDN and should be sent towards the appropriate servers.

The present invention provides numerous advantages over the prior art. According to the invention, a domain to be published to the ECDN is associated (either by static configuration or dynamically) with a set of one or more enterprise
25 zones configurable in a hierarchy. When a request for content is received from a given client in the enterprise (or otherwise accessible through the firewall), a name query to the enterprise's legacy DNS is directed (preferably via CNAME) to a set of one or more ECDN zone servers. A given ECDN zone server associates a name query with an ECDN intelligent node (or to a given content server). This

association may be based on given conditions within the network, server load conditions, or some combination of various known metrics. With ECDN virtual zones, maps that associate a given request to a server may be local or global, static or dynamic. If multiple zones are used in a hierarchy, a first level ECDN zone server redirects the name query to a second level ECDN zone server, and so on, until the appropriate response is generated. ECDN zone servers may reside on the same machine or with existing DNS infrastructure or other machines in the enterprise. This enables the present invention to be readily integrated into an existing or legacy DNS solution with minimal reconfiguration.

As described and illustrated above, a given zone may have a sub-zone and, according to the invention, that sub-zone may be managed by the enterprise or even by a third party such as an Internet content delivery network (ICDN) service provider. As an example of the latter scenario, assume that the enterprise outsources its streaming video delivery to the ICDN service provider. To this end, a host lookup for, say, allhands.streamingvideo.ecdn.company.com may be configured either statically or dynamically to a zone managed by a CDN overlay network. This has the effect of transparently allowing delivery to be seamlessly tied internally and externally via the same mechanism.

Figure 8 illustrates another example of how an enterprise may hand off management of a virtual zone to another entity, in this case an enterprise business partner. In this example, producer.com is the enterprise and consumer.com is the enterprise business partner having an end user. Producer.com desires to make given content available to the consumer.com end user, however, the producer.com namespace on which the content is published is not resolvable through public DNS and the document is not available publicly. Both producer.com and consumer.com are assumed to have implemented an ECDN system behind each of their respective firewalls, and each entity has a private namespace that they respectively manage. The goal is to make producer.com's namespace selectively available to an end user in consumer.com without, at the same time, exposing

producer.com's namespace to the public DNS. Each ECDN deploys a relay machine which, preferably, is a combination of a surrogate origin server and an ECDN zone server. Separate machines/processes may be used, or the processes may be integrated or included with other functionality. Preferably, each relay

5 machine is located within the enterprise DMZ, as illustrated in the drawing. The relay machine inside consumer.com does not know how to resolve a virtual zone associated with producer.com, but that machine does know that the relay machine inside producer.com's firewall does know how to make such a resolution. Likewise, the relay machine inside producer.com does not know how to resolve a

10 virtual zone associated with consumer.com, but that machine does know that relay machine inside consumer.com's firewall make know how to make such a resolution. As is well-known, boxes within a DMZ can be selectively allowed to communicate with one another while filtering other potentially hostile communications.

15 Referring back to Figure 8, assume that the document in question is referenced within consumer.com by a836.asp.producer.com/object.html. In this example, the zone "asp.producer.com" has been designated within consumer.com as being owned or managed by relay.consumer.com. Thus, relay.consumer.com will act as a surrogate for that document. When the consumer.com relay gets the

20 host header a836.asp.producer.com, the relay machine uses that information to connect to the relay machine in producer.com. The relay machine in producer.com gets the HTTP request and it is configured to rewrite any host header (associated with a836.asp.producer.com in this example) to foo.producer.com. The relay machine at relay.producer.com then acts as a

25 surrogate origin server for the web server at foo.producer.com, which hosts the document. The web server then returns the document via the surrogate origin server relay.producer.com, and that server rewrites the response header to point back to itself. The relay machine in producer.com then returns the data to the relay machine in consumer.com, which performs similar surrogate processing and

returns the document to the requesting end user. This completes the processing. Thus, one of ordinary skill in the art will appreciate that, in this manner, producer.com has handed off management of the relay.producer.com virtual zone to consumer.com. This arrangement creates a VPN-like connection between the
5 two namespaces to facilitate the document sharing. As illustrated, preferably the communications between the two relay machines are made over secure link using known techniques.

Having thus described my invention, the following sets forth what I now claim.

10

CLAIMS

1. A content delivery method operative in an enterprise, comprising:
establishing an enterprise content delivery network within the enterprise,
wherein the enterprise content delivery network has associated therewith a set of
5 one or more zones organized into a hierarchy; and
in response to a name query for a hostname known to be associated with
the given content being managed by the enterprise content delivery network,
returning a response to the name query;
wherein the response selectively (a) identifies a given node in the
10 enterprise content delivery network that can obtain the given content from a server
within the enterprise, (b) causes a referral to a zone within the hierarchy, or (c)
relays the name query to another path.
2. The content delivery method as described in Claim 1 wherein, if
15 the response identifies the given node, having the given node obtain the given
content from the server.
3. The content delivery method as described in Claim 1 wherein, if
the response causes a zone referral, attempting to resolve the name query at a next
20 level within the hierarchy.
4. The content delivery method as described in Claim 1 wherein, if
the response is relayed to another path, attempting to resolve the name query along
that path.
25
5. The content delivery method as described in Claim 1 further
including the step
publishing the given content into the enterprise content delivery network.

6. The content delivery method as described in Claim 5 wherein the step of publishing comprises:

identifying the hostname to be published to the enterprise content delivery network;

5 aliasing the hostname to a given zone in the hierarchy.

7. The content delivery method as described in Claim 6 wherein the aliasing step CNAMEs the hostname to a domain in which the virtual zone is prepended to at least a portion of the hostname.

10

8. The content delivery method as described in Claim 7 wherein the CNAME is generated statically.

9. The content delivery method as described in Claim 7 wherein the
15 CNAME is generated dynamically.


10. The content delivery method as described in Claim 1 wherein the name query is resolved to the IP address of the given node as a function of given network or server metrics.

20

11. The content delivery method as described in Claim 10 wherein the IP address of the given node is generated dynamically.

12. The content delivery method as described in Claim 1 wherein the
25 enterprise content delivery network is part of an Internet content delivery network (ICDN).

13. The content delivery method as described in Claim 11 wherein the enterprise content delivery network is managed as part of the ICDN.



2/4

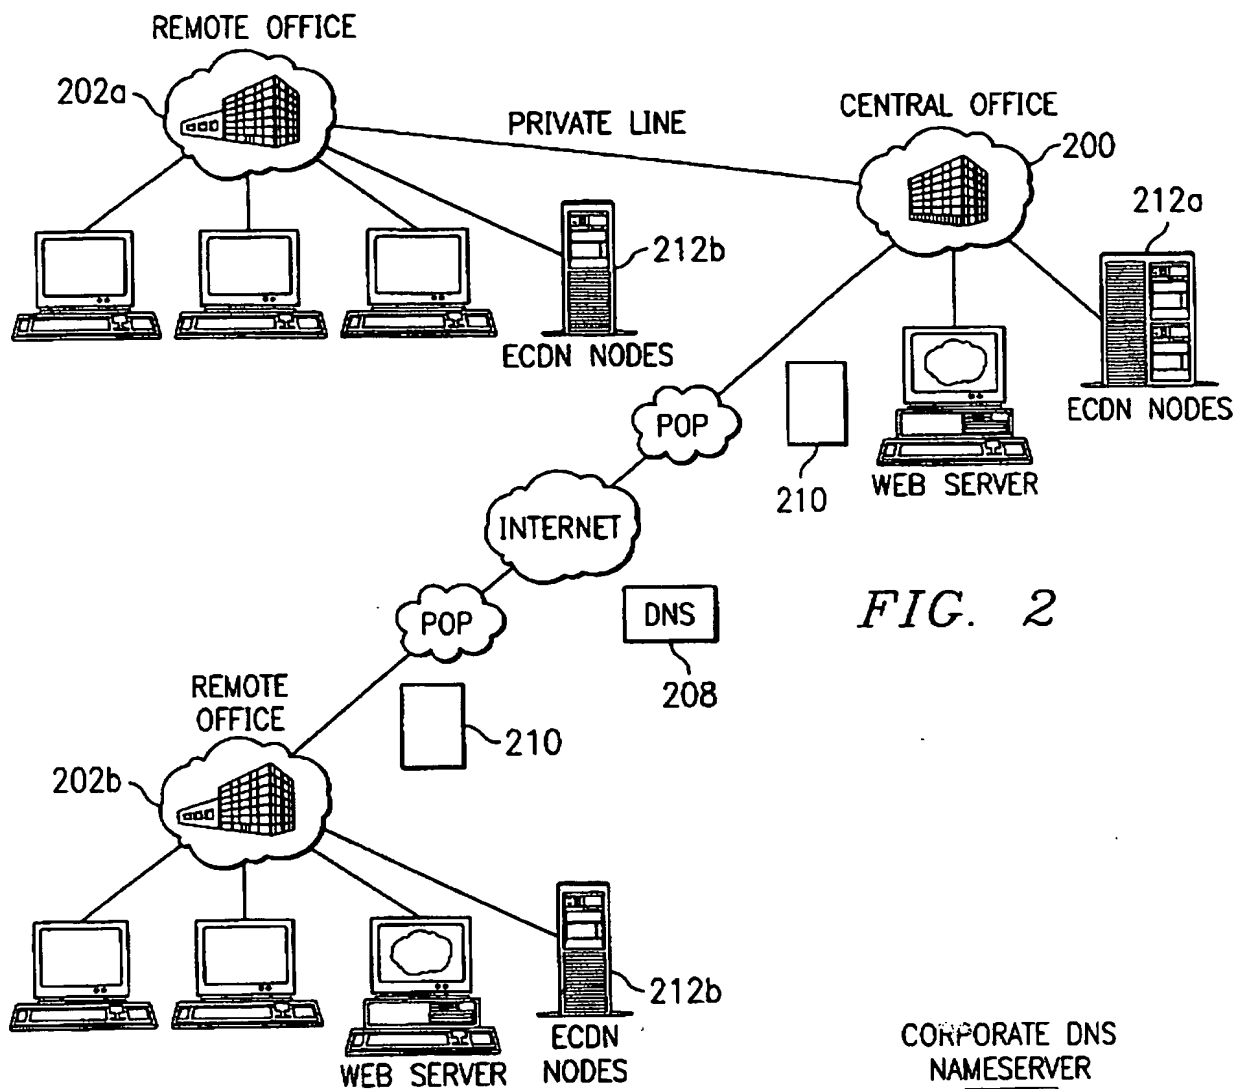


FIG. 2

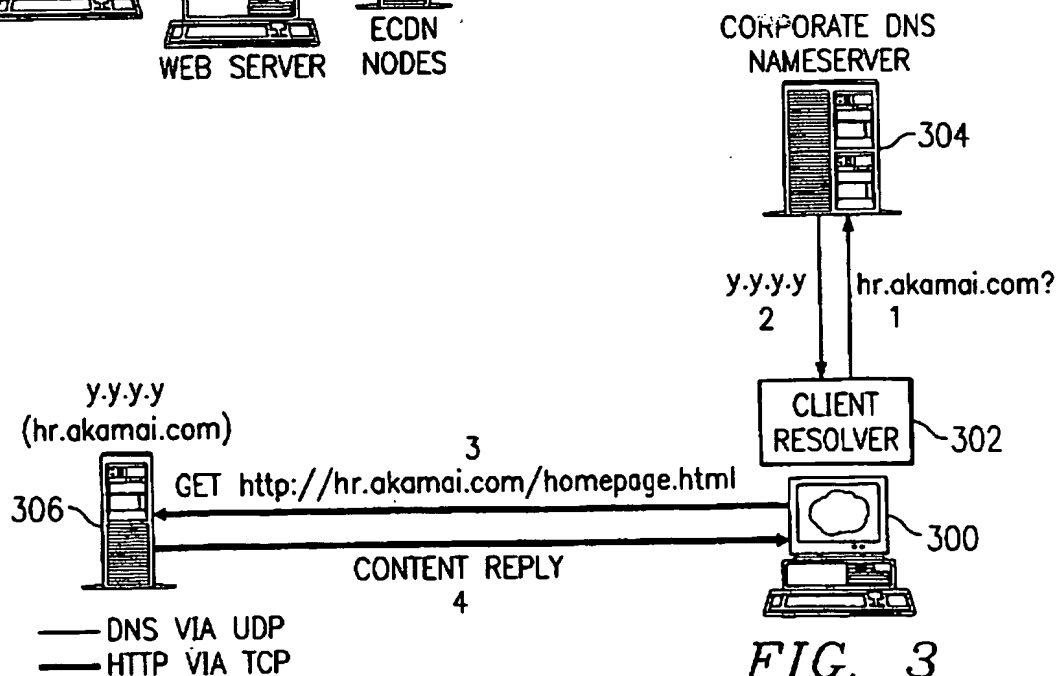


FIG. 3

3/4

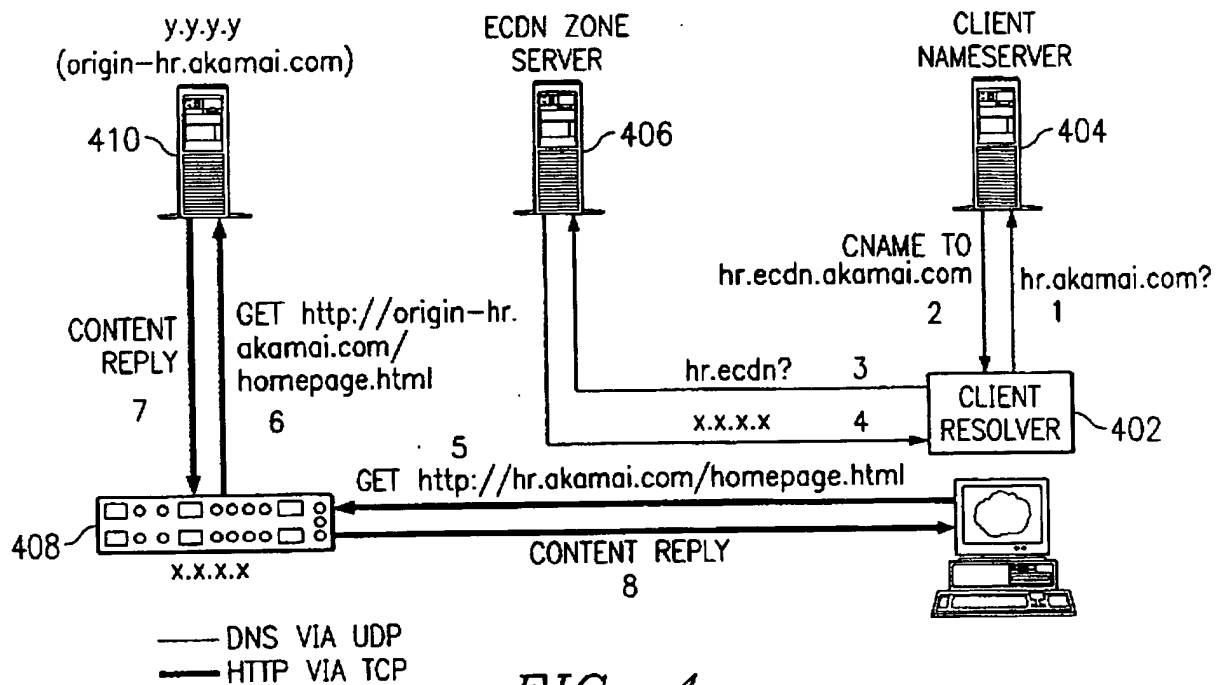


FIG. 4

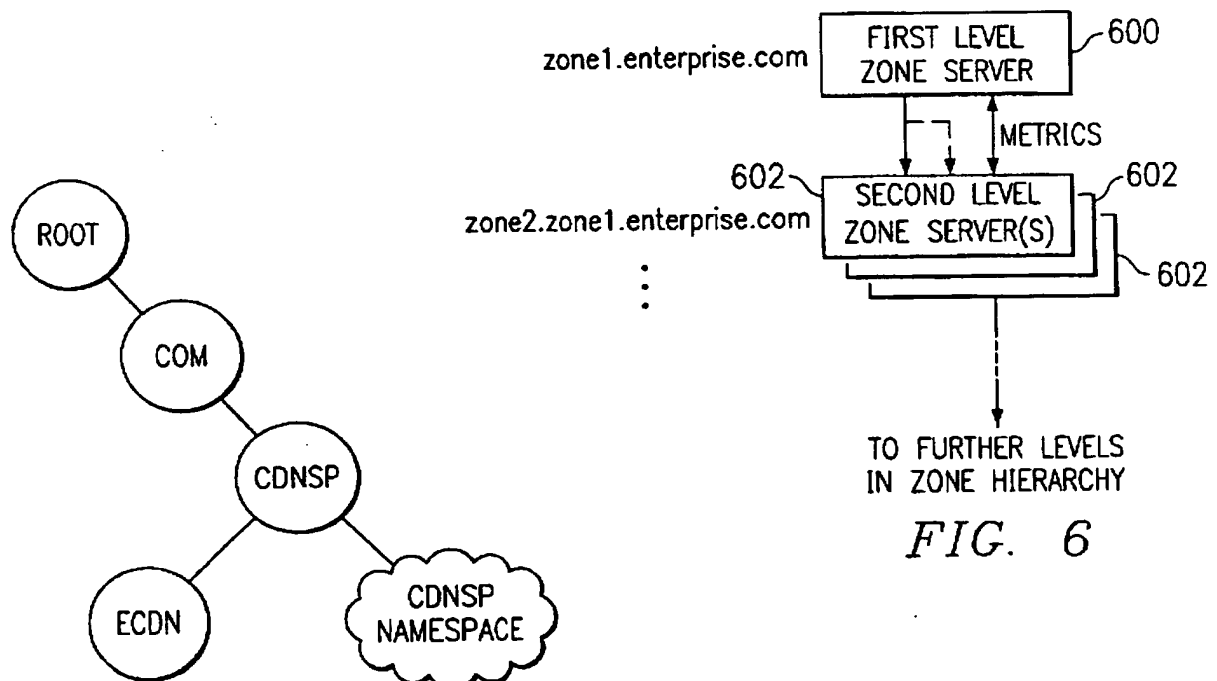


FIG. 5

FIG. 6

FIG. 7

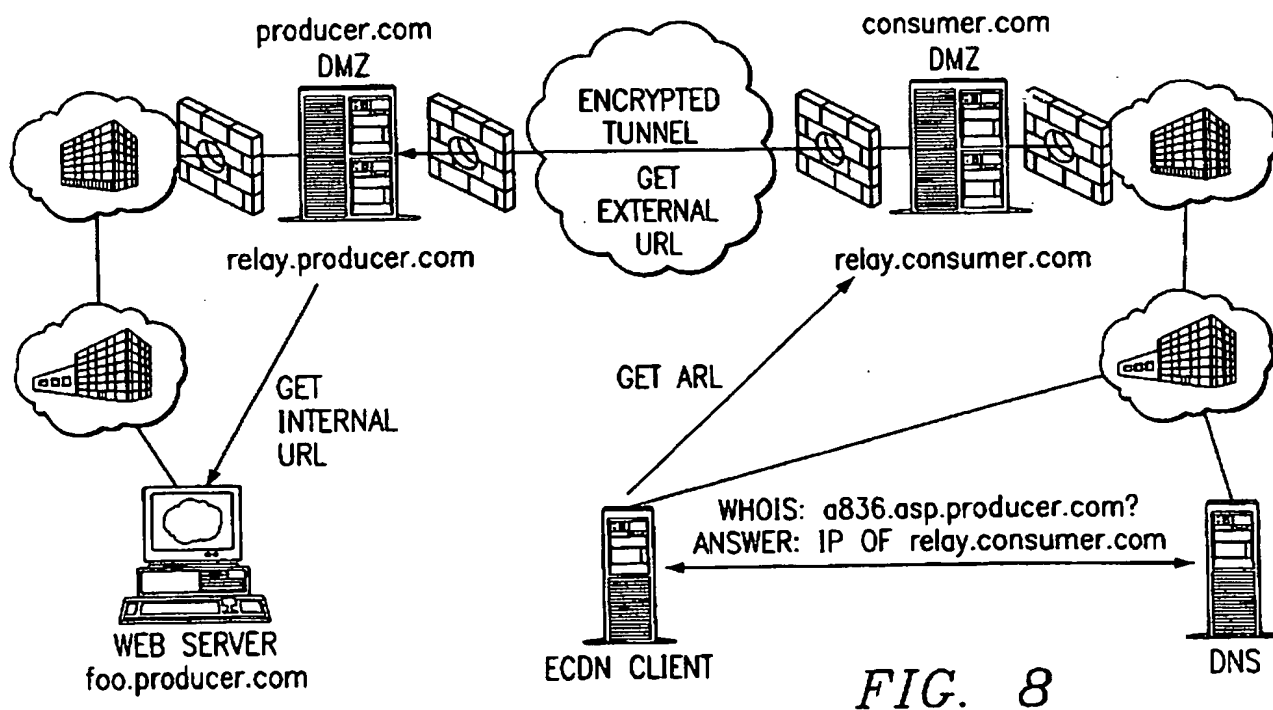
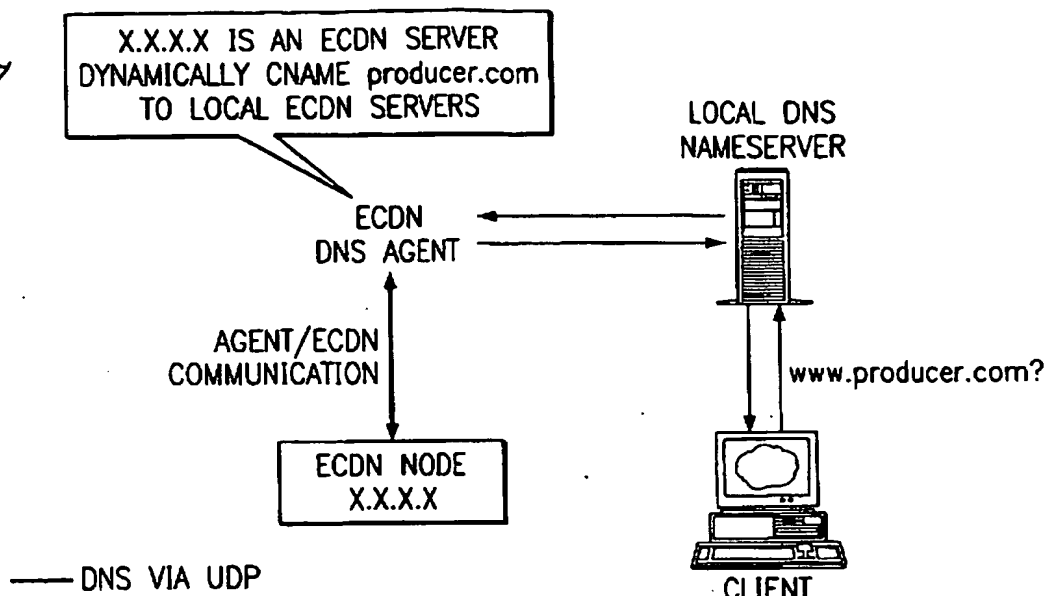


FIG. 8

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 02/01079

A. CLASSIFICATION OF SUBJECT MATTER
 IPC 7 H04L29/12 H04L29/06

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 99 40514 A (BALTER JAMES A ;FARBER DAVID A (US); GREER RICHARD E (US); SWART A) 12 August 1999 (1999-08-12) abstract page 2, line 18 -page 3, line 17 page 4, line 1 -page 5, line 24 page 6, line 24 -page 8, line 12 page 12, line 9 -page 12, line 15 page 13, line 3 -page 13, line 13 ---	1-4, 10-13
X	US 6 108 703 A (LEWIN DANIEL M ET AL) 22 August 2000 (2000-08-22) abstract column 3, line 24 -column 4, line 22 --- -/--	1-4, 10-13

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents:

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principles or theory underlying the invention

- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *G* document member of the same patent family

Date of the actual completion of the international search

5 September 2002

Date of mailing of the international search report

13/09/2002

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
 NL - 2280 HV Rijswijk
 Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
 Fax: (+31-70) 340-3016

Authorized officer

Goller, W

INTERNATIONAL SEARCH REPORT

Patent Application No
PCT/US 02/01079

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 0 817 444 A (SUN MICROSYSTEMS INC) 7 January 1998 (1998-01-07) abstract column 2, line 20 -column 3, line 41 column 5, line 50 -column 6, line 57 column 7, line 30 -column 7, line 40 -----	1,2, 10-13
X	WO 00 74347 A (ENTERA INC) 7 December 2000 (2000-12-07) abstract page 5, line 4 -page 6, line 18 page 9, line 2 -page 9, line 24 -----	1-4, 10-13

INTERNATIONAL SEARCH REPORT

Information on patent family members

al Application No

PCT/US 02/01079

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
WO 9940514	A	12-08-1999	US 6185598 B1	06-02-2001
			AU 2652999 A	23-08-1999
			CA 2320261 A1	12-08-1999
			EP 1143337 A1	10-10-2001
			EP 1053524 A1	22-11-2000
			JP 2002503001 T	29-01-2002
			NO 20004010 A	10-10-2000
			US 2002049857 A1	25-04-2002
			US 2002099850 A1	25-07-2002
			WO 9940514 A1	12-08-1999
			US 2001056500 A1	27-12-2001
US 6108703	A	22-08-2000	AU 4995299 A	07-02-2000
			BR 9912001 A	04-12-2001
			CN 1312923 T	12-09-2001
			DE 1125219 T1	23-05-2002
			EP 1125219 A1	22-08-2001
			JP 2002520735 T	09-07-2002
			WO 0004458 A1	27-01-2000
			US 2002099850 A1	25-07-2002
EP 0817444	A	07-01-1998	US 6154777 A	28-11-2000
			EP 0817444 A2	07-01-1998
			JP 10126445 A	15-05-1998
WO 0074347	A	07-12-2000	AU 4661600 A	18-12-2000
			WO 0074347 A1	07-12-2000

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ BLACK BORDERS
- ☐ IMAGE CUT OFF AT TOP, BOTTOM OR SIDES
- ☒ FADED TEXT OR DRAWING
- ☐ BLURRED OR ILLEGIBLE TEXT OR DRAWING
- ☐ SKEWED/SLANTED IMAGES
- ☐ COLOR OR BLACK AND WHITE PHOTOGRAPHS
- ☒ GRAY SCALE DOCUMENTS
- ☐ LINES OR MARKS ON ORIGINAL DOCUMENT
- ☐ REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY
- ☐ OTHER: _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.

THIS PAGE BLANK (USPTO)

